

A SOFTWARE DEVELOPMENT APPROACH
FOR PARALLEL PROCESSING SYSTEMS
INTEGRATING OBJECT ORIENTED AND FUNCTIONAL PARADIGMS

BY

BOO BRYAN BAE

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL OF
THE UNIVERSITY OF FLORIDA IN
PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF
DOCTOR OF PHILOSOPHY
UNIVERSITY OF FLORIDA

2002

To my parents, wife and daughter

ACKNOWLEDGMENTS

I would like to express my gratitude to my advisor Professor Stephen S. Yin for his advice and support in my Ph.D. program at the University of Florida and Northwestern University. Throughout my graduate study, he made contributions in various ways: advice, insight, alternative solutions and encouragement. I would like to thank my committee members, Professor Yuesi Huang, Lee, Professor Paul W. Chue, Professor Paul A. Polivinski and Professor Tim Davis. Professor Lee helped me continue my work by not only giving technical advice but also providing spiritual support. Professor Chue also provided me with kindly advice and encouragement. The discussions with him regarding parallel processing have always been productive for my study. Advice from Dr. Xiaoping Liu was invaluable to my research. I have enjoyed working with all of them. Special thanks are due to my colleagues in our software engineering research group, G. Ok, M. Chakraborty, K. Yoon, W. Song, V. Suresh, R. Yang and G. Peon. Informal discussions with them also contributed to this work. I also would like to acknowledge the partial support from Evans Laboratory, U. S. Air Force Systems Command for this research.

I wish to thank my parents for their endless support for the long years of my school life. I also wish to thank my wife, Bao Young, for her continuous encouragement and presence. Without her help, I would not be able to complete my Ph.D. degree.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iii
LIST OF TABLES	v
LIST OF FIGURES	vi
LIST OF NOTATIONS	ix
ABSTRACT	xi
CHAPTERS	
1 INTRODUCTION	1
1.1 Management of Parallelism	1
1.2 Parallel Computers	2
1.3 Distribution-Driven	4
2 COMPUTATION MODELS AND PARALLEL PROGRAMMING	8
2.1 Existing Computation Models	8
2.1.1 Imperative Computation Model	8
2.1.2 Logic Computation Model	10
2.1.3 Functional Computation Model	10
2.1.4 Object-Oriented Computation Model	11
2.1.5 Knowledge-Based Model	12
2.2 Overview of Parallel Programming Approaches	13
2.3 Parallel Language Approaches	13
2.4 Characteristics of the Dissertation	14
3 PROOF-COMPUTATION MODEL	19
3.1 Classes and Objects	21
3.2 Method Definition	26
3.3 Symbolization of Objects	26
3.4 Paramount of Objects	27
3.5 Control Functions	28
3.6 ProofPlan in PROOF	27
3.6.1 Sources of Parallelism	27
3.6.2 The Parallel Semantics of PROOF	28

4. INTERMEDIATE PROGRAM REPRESENTATION FOR PROOFL	58
4.1 Characteristics of the IPR	58
4.2 IPR Definition	58
4.2.1 Object-level IPR	58
4.2.2 Meta-level IPR	59
4.3 Transformation Rules	60
4.3.1 Representation of Interactions among Objects	61
4.3.2 Forward Transformation	61
4.3.3 Consistency Preserving Transformation	62
4.3.4 Back-rod Transformation	63
4.3.5 Examples	65
5. ALLOCATION OF PROOFL PROGRAMS	67
5.1 Task Allocation Approach	68
5.2 Object Partitioning	68
5.2.1 Modeling	68
5.2.2 Clustering at the Object level	69
5.3 Meta-level Partitioning	69
5.3.1 Graph Set Analysis	69
5.3.2 Pipelined Partitioning	70
5.3.3 Tree Partitioning	70
5.3.4 Graph Partitioning	70
6. DISCUSSION	77
APPENDIX	78
REFERENCES	79
BIOGRAPHICAL SKETCH	82

LIST OF TABLES

3.1. Multivariate testing conditions	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223	224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270	271	272	273	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288	289	290	291	292	293	294	295	296	297	298	299	300	301	302	303	304	305	306	307	308	309	310	311	312	313	314	315	316	317	318	319	320	321	322	323	324	325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340	341	342	343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360	361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378	379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396	397	398	399	400	401	402	403	404	405	406	407	408	409	410	411	412	413	414	415	416	417	418	419	420	421	422	423	424	425	426	427	428	429	430	431	432	433	434	435	436	437	438	439	440	441	442	443	444	445	446	447	448	449	450	451	452	453	454	455	456	457	458	459	460	461	462	463	464	465	466	467	468	469	470	471	472	473	474	475	476	477	478	479	480	481	482	483	484	485	486	487	488	489	490	491	492	493	494	495	496	497	498	499	500	501	502	503	504	505	506	507	508	509	510	511	512	513	514	515	516	517	518	519	520	521	522	523	524	525	526	527	528	529	530	531	532	533	534	535	536	537	538	539	540	541	542	543	544	545	546	547	548	549	550	551	552	553	554	555	556	557	558	559	560	561	562	563	564	565	566	567	568	569	570	571	572	573	574	575	576	577	578	579	580	581	582	583	584	585	586	587	588	589	590	591	592	593	594	595	596	597	598	599	600	601	602	603	604	605	606	607	608	609	610	611	612	613	614	615	616	617	618	619	620	621	622	623	624	625	626	627	628	629	630	631	632	633	634	635	636	637	638	639	640	641	642	643	644	645	646	647	648	649	650	651	652	653	654	655	656	657	658	659	660	661	662	663	664	665	666	667	668	669	670	671	672	673	674	675	676	677	678	679	680	681	682	683	684	685	686	687	688	689	690	691	692	693	694	695	696	697	698	699	700	701	702	703	704	705	706	707	708	709	710	711	712	713	714	715	716	717	718	719	720	721	722	723	724	725	726	727	728	729	730	731	732	733	734	735	736	737	738	739	740	741	742	743	744	745	746	747	748	749	750	751	752	753	754	755	756	757	758	759	760	761	762	763	764	765	766	767	768	769	770	771	772	773	774	775	776	777	778	779	780	781	782	783	784	785	786	787	788	789	790	791	792	793	794	795	796	797	798	799	800	801	802	803	804	805	806	807	808	809	810	811	812	813	814	815	816	817	818	819	820	821	822	823	824	825	826	827	828	829	830	831	832	833	834	835	836	837	838	839	840	841	842	843	844	845	846	847	848	849	850	851	852	853	854	855	856	857	858	859	860	861	862	863	864	865	866	867	868	869	870	871	872	873	874	875	876	877	878	879	880	881	882	883	884	885	886	887	888	889	890	891	892	893	894	895	896	897	898	899	900	901	902	903	904	905	906	907	908	909	910	911	912	913	914	915	916	917	918	919	920	921	922	923	924	925	926	927	928	929	930	931	932	933	934	935	936	937	938	939	940	941	942	943	944	945	946	947	948	949	950	951	952	953	954	955	956	957	958	959	960	961	962	963	964	965	966	967	968	969	970	971	972	973	974	975	976	977	978	979	980	981	982	983	984	985	986	987	988	989	990	991	992	993	994	995	996	997	998	999	1000
--------------------------------------	---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	------

LIST OF FIGURES

3.1	The interface of the class <code>BoundedBuffer</code> $\langle \text{class } B \text{ of } \text{Buffer} \rangle$	23
3.2	The definition of the class <code>BoundedBuffer</code> without guard constraints	23
3.3	The definition of the class <code>ExtendedBuffer</code> $\langle \text{class } E \text{ of } \text{Buffer} \rangle$	24
3.4	Interface of the classes <code>ConsumerClass</code> and <code>ProducerClass</code>	25
3.5	The set of the objects for a producer-consumer problem $\langle \text{set } S \rangle$	25
3.6	An array partitioning problem $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	26
3.7	A solution for the array partitioning $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	26
3.8	The definition of the class <code>BoundedBuffer</code> $\langle \text{class } B \text{ of } \text{Buffer} \rangle$	31
3.9	The complete definition of the class <code>ExtendedBuffer</code> $\langle \text{class } E \text{ of } \text{Buffer} \rangle$	32
3.10	The definition of the objects for producer-consumer problem $\langle \text{set } S \rangle$	33
4.1	Two-level transformation $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	38
4.2	Pseudocode relations among functions $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	38
4.3	The Petri-net representation of (a) value flow (b) shared flow	39
4.4	The semantics of the function application $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	42
4.5	The operational semantics of the selector node $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	43
4.6	The operational semantics of the distribution node $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	43
4.7	The operational semantics of the list handling nodes (a) construct node (b) split node (c) merge node	44
4.8	Petri net representation of <code>nd</code> $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	47
4.9	The semantics of the sequential iteration $\langle \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \text{ of } \text{array } A \rangle$	49

4.10	The semantics of the concurrent operations	68
4.11	The semantics of the selective operations	70
4.12	The semantics of the mutually exclusive access	71
4.13	The semantics of the communication between two objects	72
4.14	The semantics of the method invocation with a guard	74
4.15	The transformation rules for function and functional forms	75
4.16	The transformation rules for the nested functions	76
5.1	The modeling of the object behavior	80
5.2	An object clustering for the flow during photographer's problem	83
5.3	An LPT representation for M.	102
5.4	Sample gas flow examples	112
5.5	A task precedence tree	113
5.6	A schedule obtained from McCreary approach	114
5.7	A gas tree for the tree precedence example	115
5.8	A schedule obtained from our approach	116
5.9	A task precedence graph for the FFT problem	120
5.10	A gas graph for the FFT problem	122
5.11	A schedule for the FFT problem	123

LIST OF NOTATIONS

Symbol	Description
A	a labeled transition system
$A_0(P)$	an initializing assertion of the program P
$\alpha(x,y)$	a communication line between node x and y
B	a set of final places in the Petri net
$c(a)$	an execution time required to complete a node a
E	a set of arcs in the method level IPL
$E(x,a)$	a execution time to complete a given task x with a processors
$E(x_i)$	one-pass completion time for segment a_i in a pipelined process
F	a set of arcs representing flow relations in the Petri net
G	a directed graph representing the method level IPL
$G(IPN)(a_1, \dots, b)$	a gain obtained by clustering nodes a_1, \dots, b
G_p	a gain graph
G_t	a task precedence graph
I	a set of initial places in the Petri net
N	a general set
P	a set of places in the Petri net
pp	a port of a place p
pt	a port of a place p
PN'	a Petri net used for the object level IPL
R	a predicate function to modify the object state
tr	a port of a transition t
tr	a port of a transition t

\mathcal{T}	a set of transitions in the Petri net.
\mathcal{T}_0	a place term
\mathcal{T}_n	a basic polynomial term.
\mathcal{V}	a set of nodes in the method level IFG.

*Abstract of Dissertation Presented to the Graduate School of
the University of Florida in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy*

**A SOFTWARE DEVELOPMENT APPROACH
FOR PARALLEL PROCESSING SYSTEMS
INTEGRATING OBJECT-ORIENTED AND FUNCTIONAL PARADIGMS**

by
Don-Eun Kim

May 1992

Chairman: Stephen S. Yim
Major Department: Computer and Information Sciences

Although there are many parallel computers available, the effective use of their computing power is difficult to achieve due to the lack of proper software development methods and techniques for parallel processing systems. In this dissertation, we present a computation model, called PFOOP (Parallel Object Oriented Functional) as a basis for expressing parallelism, and an intermediate form, called IFB (Intermediate Program Representation) as a basis for exploring parallelism in software development for parallel processing systems. We also present graph-like transformation techniques on different types of parallelism. We develop the computation model PFOOP in which the object-oriented paradigm is integrated with the functional paradigm without sacrificing the advantages of either. The major features of our computation model include expressing the various granularity levels of parallelism, integrating the referential transparency and history sensitivity, and supporting abstraction and specialization without interference. The computation model PFOOP serves as a basis for developing our approach to the software development for parallel processing systems. We develop the IFB, which is used for representing all the parallelism in the software and performing partitioning analysis. We develop

transformation rules from a PROOF code to the IPR and then the IPR to a target code and show that the transformation rules preserve the semantics of the original PROOF program. For efficient exploration of parallelism, we present the grain size determination algorithms on various patterns of parallelism and compare them with the existing approaches. Software development based on PROOF has the following distinct advantages. First, parallel aspects of the software can be coded as a grain size, using object oriented concepts. Second, our approach can reduce the communication and synchronization costs by automatically subdividing proper codes during the transformation steps. Third, the use of the IPR improves portability of the software by separating architecture-dependent codes from the semantics of the software.

CHAPTER I INTRODUCTION

With the advances in computer technology, the speed of computers has significantly increased over the last several decades. However, even such rapidly increased speed is not sufficient to hold our doors to solve complex problems, involving neural and intelligent, robotics, high energy physics, molecular physics and space sciences. Parallel processing appears to be a promising solution to satisfy such problems. Unfortunately, software development methods and techniques for parallel processing systems are far from being mature enough to effectively utilize their potential. The lack of such methods and techniques is a major obstacle to the widespread use of parallel processors in a variety of applications areas. The goal of this research is to develop a software development approach for parallel processing systems and study the issues involved in such development.

Software development is more complex for parallel processors than for sequential computers in at least two aspects: program correctness and efficiency. One of the requirements for the correctness of the parallel programs is the independence of the execution results on the number and speed of processors executing the program. This requirement can be met only if the parallel tasks are properly synchronized. The focus on efficiency in parallel program design increases the complexity of software development. The efficiency of parallel implementation is usually measured by dividing the speed-up of the parallel program relative to its sequential version with

the number of processors involved is that parallel program execution. Thus, in order to achieve a maximum performance in terms of efficiency, proper techniques for software development for parallel processing systems need to be developed.

1.1 Management of Parallelism

Parallel processing of a program is based on the following two assumptions:

- the availability of the parallel computers that can cooperate to solve a given problem,
- the existence of program development techniques such as design techniques for parallel programs and efficient implementation strategies.

In fact, as cost effective parallel computers become more available on the market, greater computing power can be easily obtained. Thus, the first assumption seems to be met. However, current software development techniques for parallel processing systems are not adequate to specifically address the issues involved in programming parallel processes, most importantly parallelism.

The essence of parallel processing is to manage parallelism. Parallelism refers to the execution of some portions of a program in a simultaneous manner. The management of the parallelism involves two issues:

- Expressing parallelism: for parallel processing of a program, the parallelism in that program needs to be represented using a programming language, and
- Exploiting parallelism: the parallelism needs to be efficiently exploited on the underlying parallel computers.

Parallelism can be specified either implicitly or explicitly. Explicit parallelism is expressed using explicit parallel constructs by the programmers. Thus, the detection of the explicit parallelism is easy. However, the programmers have to ensure

the correct use of such constructs in terms of communication and synchronization of parallel programs. In contrast, *implicit* parallelism implies that programmers are not concerned about the parallelism. In that case, the underlying computation models inherently support the parallelism. Thus, programmers do not use any explicit parallel constructs. On the other hand, either a compiler or run time system requires to detect parallelism in the programs. Once the parallelism is explicitly expressed and detected, it needs to be exploited on parallel computers. The exploitation techniques vary depending on the kind of parallelism and the parallel computers involved. Exploiting and exploiting parallelism in the software is most important part of software development for parallel processing systems and is the subject of this study.

1.2. Parallel Computers

In this section, we briefly summarize the existing parallel computers. These can be divided into two categories: *special-purpose* and *general-purpose*. The *special-purpose* parallel computers are designed to solve specific problems. The vector processors in super computers, such as Fujitsu VP 300 and Hitachi S-81P, have been developed to exploit parallelism mainly within loops. In these computers, the vector processors are controlled by a main control unit. Processor arrays consist of a number of identical processors under the control of a common control unit. Each processor has access to its own data, and thus the same operation can be performed simultaneously on many data items. In Flynn's classification [28], this kind of computer belongs to SIMD (Single Instruction Stream, Multiple Data Stream). The parallel computers belonging to this category include Connection Machines [29], DAP (Data Driven Array Processor) [41] and IBM GF11 [11]. *System arrays* [32] consist of

¹These computers can be classified as general-purpose computers. However, because they support parallel processing a limited only to vector processing, we classify them as special-purpose computers.

ten or more dimensions of array processing in which each processor has its own unit to perform specific task. The data stream passes through the array processors to produce a final result.

In comparison to these special purpose parallel computers, the general purpose computers belong to MIMD (Multiple Instruction Stream, Multiple Data Stream) according to Flynn's classification. The parallel computers in this category are a set of asynchronous parallel processors. These computers can be further divided into shared memory parallel computers and distributed memory computers. Because the focus of this research is to develop a parallel programming approach for general-purpose parallel processors, we will further examine each of the two kinds of parallel computers.

An MIMD shared-memory parallel computer consists of a number of processors all having access to a single shared memory. The processors communicate by read and write operations, and are connected to the shared memory via one or more shared buses or interconnecting networks. As the number of processors in the system increases, the communication overhead becomes a bottleneck in terms of performance as well as cost. Thus, a linear speedup with an increase in the number of processors is not achievable or is limited to a certain number of processors. In the MIMD shared-memory parallel computers, the major software design problems include data access synchronization and load balancing. The shared memory parallel computers include Cray X MP and Y MP series [4], Alpha FPA [11], Encore Multimax [65]-[68] MP3 [66] and Sequoia Redstone [77].

To model the computations on the shared memory parallel computers, a theoretical model of parallel computation, called PRAM (Parallel Random Access Memory) [69] can be used. In the PRAM model, each processor is a RAM (Random Access Machine). Processors communicate by reading from and writing to a global memory

The instructions in each processor are executed in a synchronous manner as if there were a global clock. PRAM models can be further divided according to the access constraints for the simultaneous read and write access. Exclusive Read Exclusive Write (EREW), Concurrent Read Exclusive Write (CREW), Exclusive Read Concurrent Write (RCRW) and Concurrent Read Concurrent Write (CRCW). Although the underlying assumptions in such PRAM models, such as constant communication time and synchronous execution of instructions, are not realistic, they have been used as the multiprocessor representations in the evaluation of parallel algorithms and task scheduling problems.

An MIMD distributed parallel computer consists of a set of processors each having a non shared local memory. Processors communicate by message passing via communication channels. In the MIMD distributed memory parallel computers, the synchronization is implicit through communication. The popular interconnection networks include hypercube, ring, tree and mesh. Unlike shared memory parallel computers, distributed memory parallel computers are not affected by the memory contention problem and are more easily expandable. One of the problems encountered by the distributed memory parallel computers is a message passing latency due to the possible transferring of data via intermediate processors. The major software design problems include data placement, communication overhead and scheduling. The shared memory parallel computers can be seen as a special form of distributed memory parallel computers in which all the processors are fully connected. Our software development approach is targeted for the MIMD distributed memory parallel computers. Any approach for the MIMD distributed memory parallel computers can be adopted to MIMD shared memory parallel computers with minor modifications. The MIMD distributed memory parallel computers include hypercube [26], SCUBC [24], DDM [24] and Janus Transputer network [24].

1.3. Dissertation Outline

In this dissertation, a software development approach for parallel processing systems is presented. The approach is based on an integrated object-oriented and functional paradigm. We develop a computation model PROOF (Parallel Object-Oriented Functional) based on object-oriented and functional paradigms [84]. This PROOF computation model permits the expression of explicit and implicit parallelism using objects and functions at various levels of granularity. This computation model serves as a platform in developing our approach to programming parallel processors. It provides a framework for programmers to express parallelism in the application. Once parallelism is expressed using PROOF, it can be exploited when the PROOF program is transformed to a target code. The transformation is done in two steps. The first step transforms the PROOF program to the IFN which is independent of both a target language and a target machine. We do analysis on the IFN to determine the proper grain size. The second step transforms the IFN to a target code, using the information generated by the grain size analysis. The dissertation is organized as follows:

Chapter 2 surveys computation models and software development approaches for parallel processing systems². Three parallel programming approaches—parallelizing compiler approach, programming/language constructs approach and parallel programming language approach—are discussed. In particular, the existing parallel programming language approaches have been extensively discussed. Chapter 3 presents the PROOF computation model. In this chapter, the characteristics of the PROOF are presented along with examples. Major characteristics include incorporation of object-oriented concepts to functional paradigm, allowing various granularity levels

²We have abstracted the phrase “software development for parallel processing systems” to “programming parallel processors”.

of parallelism, integration of referential transparency with history sensitivity and integration of parallelism with subroutines. Chapter 4 presents the IPR and the rules for two-phase transformation of the PROOF/L² programs. A PROOF program is first transformed to the IPR code, and then the IPR code is transformed to a target code. We show that such a transformation preserves the semantics of the original program. Chapter 5 presents an allocation strategy for exploiting parallelism in programs based on the computation model PROOF. We develop grain size determination algorithms on various types of parallelism. In the case of pipeline parallelism, we show that our algorithm can find optimal solutions. In the case of tree parallelism and graph parallelism, we introduce the concept of gain to analyze grain sizes. We also show that our approach performs better than the existing approaches. Chapter 6 summarizes the dissertation and discusses the directions of future research.

²The PROOF/L is a prototype programming language based on the computation model PROOF. The syntax of the PROOF/L is given in Appendix A.6.

CHAPTER 2 COMPUTATION MODELS AND PARALLEL PROGRAMMING

In this chapter, we have surveyed the existing computation model and their supports to parallel programming, and the existing approaches to programming parallel processes. This chapter is organized as follows. In Section 2.1, the existing computation models and its extensions for parallel processing are surveyed. In Section 2.2, the existing parallel programming approaches are compared, and the parallel language approach is identified as the most promising. In Section 2.3, the existing parallel programming approaches are surveyed. In Section 2.4, the contributions of this dissertation are outlined.

2.1 Existing Computation Models

As the various cost-effective MIMD parallel computers begin to appear on the market, more general approaches to programming such computers are required so that the programmers can write programs without worries about the details of the parallel computers. Because programming languages have been used as communication vehicles between the programmers and the computers since the appearance of the first computer, the programming languages in writing parallel programs play an important role in the parallel programming systems. More importantly, the underlying computation model on which the semantics of the programming language is based has significant impact on the parallel programming systems. There are many programming languages used in writing programs for these parallel processing systems. Since each of these programming languages has different features, their comparison

is not an easy task. Rather than examining each of these programming languages in terms of the specific features, we classify them into five different categories – imperative, logic, functional, object oriented and knowledge based – according to their underlying computational model.

3.1.1 Imperative Computational Model

As Backus [6] described, this programming style is sequential in nature, the programming languages in this category have been significantly influenced by the von Neumann architecture. In imperative programming, computation is achieved by the repeated and step by step-computation of low level values and the assignment of these values to memory locations. Programmers explicitly specify the order of instruction execution using control structures. This is also called as the von Neumann model of programming.

There have been ways to use the programming languages in this model for parallel processing. In one, the programmers write sequential programs and let parallelizing compilers detect parallelism in the programs. This has been widely used primarily because the programmers need not make additional effort in coding their programs to execute in parallel. However, these compilers can only detect localized parallelism within loops.

Another way is to extend existing imperative languages with parallel language constructs. Since the parallel constructs are added to inherently sequential programming languages, the programmers are heavily burdened with explicitly specifying parallelism, synchronization and synchronization among parallel units. Fortran [3, 36], C [3, 36], Ada [32], and Occam [34] belong to this category.

2.1.2 Logic Computation Model

In logic programming, the basic element is the clause. A logic program consists of a set of clauses: goal clause, assertion clause and conditional clauses. declaring the goal, fact and properties that describe the problem for which the solution is sought. The implementation system(programter), then, uses a verification process to show if the goal clause is true or not. One goal may consist of a number of sub-goals and may be solved by solving all sub-goals independently. Another goal may be solved by solving only some of sub-goals. This gives rise to a great deal of potential parallelism: AND ϕ parallelism, OR parallelism and Disjunct-parallelism.

Efficiency of execution is a major problem in this model of programming. Thus, even if a logic programming language might be suitable for a certain application, lack of efficiency may restrict its use primarily to a rapid prototyping tool. The programming languages belonging to this category include PROLOG [36] and Concurrent PROLOG [7].

2.1.3 Functional Computation Model

In imperative programming, the basic element of programming is the assignment statement, which modifies a value of a variable as a result of computation. In functional programming, the basic element is a function, which receives input and produces output as a result of function application.

A function describes the relationship between input and output. The essence of the functional computation is to combine such functions and produce more powerful functions so that a solution for a problem can be obtained by applying such functions to input. The functional languages based on this model are referentially transparent and the programmers cannot mis-face race conditions. As a result, this model has great potential of exploiting implicit parallelism by removing side-effects caused by

assignment statements. Functional programming has been one of the main directions in developing new languages that directly address the challenge of parallel programs. In functional languages, computational abstractions are expressed through functions. A first-order function takes data objects as arguments and produces new data objects as results. The function abstracts the method used to produce the new objects from the arguments. High-order functions generalize this further by taking data objects as well as other functions as arguments and producing new data objects and new functions. It can be further advanced by the relational transparency. Because the arguments of a function could be evaluated in any order, functional programming languages promise to be more portable and easily maintained through an entire software life cycle.

However, due to its history consistency, the expressive power of programming is limited. This model is not suitable for expressing inherently concurrent nature. The data flow model shares many common properties with the functional model. In the data flow model, data "flows" from one statement to another, execution of statements is data driven, and statements obey the single assignment rule. In a data flow computer, the availability of input operands triggers the execution of the instructions which consume the inputs. The programming languages belonging to this category include Pascal [47], SIAL [49] and Modula [34].

3.1.4. Object-Oriented Programming Model

The object object programming model is based on the concept of an object. A program is a set of objects. Each object is a self contained entity with its own private data and a set of methods to manipulate those data. Any access to the data of an object must be done by calling an appropriate method. The behavior of an object is defined by its class, which is a description of a set of objects having a common

behavior. An abstract mechanism allows a class to be defined as an extension of another class.

In a traditional object-oriented programming model, there is always one active object at a time. The active object can send a message to a receiver object. Then the receiver object becomes active and the sender object must wait. After the receiver object returns a result to the sender object and becomes inactive, the sender object can continue its computation. To incorporate parallelism in this model, a number of means have been proposed [3]:

- 1) allow an object to be active without receiving messages,
- 2) allow the receiving object to continue execution after it returns its result
- 3) broadcast a message to several objects at once
- 4) allow the nodes of a message to proceed in parallel with the receiver.

Any combination of these can be used.

An object-oriented programming model naturally reveals existing parallelism in programs [32]. Besides the advantages of modularity, maintainability and reusability, one advantage of this model over others is that the concept of an object can be used at earlier stages of software development cycles than the implementation stage. It implies that parallel processing aspects such as parallelism and communication among parallel components can be naturally handled at the earlier stage of the software development/development. Consequently, it is easy for the programmers to handle parallelism and communication among parallel components. However, in this model, parallel execution among concurrent objects is the only source of parallelism, and the amount of parallelism to be exploited may not be sufficient for effectively utilizing many fine-grain processors. The programming languages belonging to this category include ABCG/1 [33], Arc 1 [34], Concurrent Smalltalk [35], Emerald [32] Interat [36] and POOL [4].

3.1.4 Knowledge-based Model

In addition to these computation models we presented, there are at least three computation models which perform their computation in different ways from the previous models. The rule-based model performs the computation by using rules, in a production system language, such as DPH [14], the basic concepts are IF-THEN statements. These rules are repeatedly executed until none of IF conditions holds. In cellular models, the computation is performed by a set of cells. Each cell performs the computation by communicating with its neighbors in a regular pattern that matches the flow of data and control in the target computation. The class of programming language belonging to this model is a message network language, such as NETL [25]. In neural models, the basic concepts are statements defining the network topology together with the recall (and learning) functions of the neurons. The execution occurs each time the network is updated. The languages belonging to this group can be regarded as knowledge-based languages and are used for applications in artificial intelligence, cognitive psychology, and learning systems.

3.2 Overview of Parallel Programming Approaches

There are three approaches to programming parallel computers. One approach is to write programs using conventional sequential programming languages, such as Fortran [1, 26, 27] or C [4, 28], and parallelize the programs using parallelizing or vectorizing compilers. Although this approach seems to be attractive since many types of existing sequential software can be adapted to such a parallel programming environment with some modifications, it is hardly an effective approach. The parallelizing or vectorizing compilers fail to control most of the parallelism. They can only detect parallelism associated with iterations over common data structures, such as arrays and matrices [7], and require extensive dependency analysis [8]. Usually, sequential

language are extended with compiler directives in order to help the compiler detect parallelism. Because these extensions are machine-specific, portability of programs is hampered. Even after extensive research and development, parallelizing compilers have not met the expectations that were raised for them. The problem is this approach is not the compiler themselves, but the inherent sequential characteristics of imperative programming languages. As we discussed above, these languages are designed for sequential execution or sequential processes. Although this approach is very popular in scientific application areas, it does not provide promising directions for the future of parallel programming.

Another approach is to use parallel language constructs to explicitly model the parallelism in programs. These parallel language constructs include the parallel statement and signal, output commands in CSP [46], master and slave, agent operations in Concurrent Pascal [33], and task and module mechanisms in Ada [35]. Although in this approach the imperative languages are extended with some language constructs, the basic model of computation is still sequential as in the first approach. Using the parallel language constructs, it is the programmers' responsibility to explicitly express the parallelism and ensure the correct communication and synchronization among parallel units, which is an extremely complex and difficult task. This is one of the major flaws in software systems for parallel computing, and there is no way solution is right. In addition, these parallel language constructs are only suitable to express coarse grain parallelism. Thus, master and slave parallelism cannot be expressed in this approach.

The third approach is to use parallel programming languages, such as M-Heyrovsky [95], and SISCAL [96], which are functional languages tailored for scientific computation, PARLOG [97], a parallel logic language, and Act-1 [98], an object-based language based on the Actor model. The underlying computation models of these parallel programming languages are fundamentally different from the underlying models of imperative programming languages in that parallelism is mostly implied and coarse-grained parallelism is easily obtainable. Hence, the programmers using these languages are liberated from the complications caused by parallelism. This is considered the most promising approach to parallel programming, and our computation model belongs to this category. We will discuss it further in the following section.

2.1. Parallel Language Approaches

Tucker [99] has identified three fundamental issues to be resolved for parallel execution of a program: 1) identifying parallelism in the program, 2) partitioning the program into sequential tasks, and 3) scheduling the tasks on processors. Although these issues have not been treated as a whole, they are closely related in the sense that the choice of a strategy in one area may significantly affect the choice of a strategy in other areas. Tucker introduced a simple task method for automatically partitioning data flow graphs. The goal is to find a schedule to minimize the completion time of the program by avoiding exploitation of top fan as a guide of parallelism. This approach encompasses three steps: 1) construct the program graph, 2) partition the graph and 3) generate the target code. In step 1), the execution time is assigned to each node and the communication time between two neighboring nodes is also assigned to the edge connecting the two nodes. In step 2), the program graphs are partitioned. Finally, in step 3), the macro-data flow modules are generated. Each macro data flow module is a code segment to be executed sequentially in one processor. This method

has been applied to the SP44 programs, and is reported as the most sophisticated methods to partition a program. However, this automatic partitioning system can only exploit parallelism appearing in the program graph. Some patterns of parallelism cannot be shown in the program graph, and thus the compiler usually does not have enough information to detect such parallelism and exploit it in a proper way. For instance, pipeline parallelism is almost impossible to exploit in this approach since the compiler itself cannot detect pipeline parallelism. In addition, the parallelism due to the arbitrary producer-consumer relationship cannot be exploited at all.

Gallberg [24] developed a method to programming parallel processors for functional programs by introducing a logical construct called a serial combinator. A serial combinator is defined as a function with the following properties: 1) its body contains no free variables; 2) its body is sequential and contains constructs for synchronizing its execution with other tasks; 3) its body could not occur as a subexpression within the body of another serial combinator. In this approach, the third property implies that the programmers have to determine as few serial combinators as possible, since they cannot be collapsed to form a bigger serial combinator. It also implies that the program developed for one parallel computer may not be directly portable due to possible performance degradation. In addition, the programmers must ensure the correct synchronization and communication among tasks.

In [25, 26], Repnstein et al. introduced an approach to programming a network of workstations based on a coarse-grain dataflow concept. Their computation model incorporates two different paradigms: at the lower level procedural language is used to express functions or procedures, at the high level a declarative style of programming is used to express implicit parallelism among the functions written in procedural programming languages. Foster and Taylor [21] introduced Stream for

parallel programming, which is based on logic computation model. Strand can provide an interface to other languages as in [47, 48]. They have ignored the issue of the granularity of parallelism and assume that the programmer will make choices on the grain size during the development of the application. Another approach based on the coarse-grain dataflow computation model was introduced by Goudeti and Lee [49]. Their approach was aimed at the use of shared memory parallel processors for solving solvable problems.

Although Linda [34] is not based on any specific computation model or approach, it is worth comparing it with other approaches. Linda is a small set of operations that can be added to a base language to create a parallel processing dialect. The concept of Linda is based on the tuple space of parallel processing. Processes and data can be considered to be elements in tuple space. Communication between processes occurs in the following way: the reader creates data in tuple space; the receiver gets the data in the tuple space, hence communication takes place. Linda provides the following four basic operations: *in*, *out*, *rfind* and *read*. *in* removes the tuple *out* from the tuple space, *rfind* reads the tuple, but leaves the tuple to be used by other processes, *out* creates a new tuple and places it in the tuple space, and *read* creates a new tuple by generating a process. A disadvantage of Linda is that the programmer has to write programs in terms of communication with other processes. In addition, its implementation on distributed memory parallel processors requires significant overhead to support communication via shared memory. Another disadvantage in Linda is that the use of conceptually shared memory does not provide any protection from erroneous access to any data item in the shared memory. When we consider the fact that data or messages need not be accessible to any other process than those processes which used them, Linda cannot support an information hiding principle.

In conclusion, these existing approaches share some or all of the following disadvantages:

- Software engineering concepts for managing parallelism have not been fully incorporated.
- Most approaches are targeted for the shared memory processors.
- The concept of shared data has not been introduced into programming parallel processors.
- Coding of correct synchronization and communication using explicit constructs is still the programmer's responsibility.

In order to develop software for parallel processing systems, much effort is still needed to address these problems.

2.4. Contributions of the Dissertation

In developing an approach to software development for parallel processing systems integrating object-oriented and functional paradigms, the dissertation makes the following contributions:

- The Parallel Object-Oriented Functional Computation Model (POOFCM) has been developed. In POOFCM, the object-oriented concepts have been integrated to the functional paradigm without sacrificing the advantages of either paradigm. POOFCM offers the following advantages: 1) Parallelism is different granularity level can be exploited. 2) By separating synchronization mechanisms from the operation of the method, inheritance and parallelism can be integrated without interference. 3) Parallel aspects of the software can be treated as a primary issue at the early stage of the software development so that the complexities at the later phases of the software life cycle can be reduced.

- The intermediate form, called *Intermediate Program Representation (IPR)*, for the internal representation of the PROOF/L programs has been developed. The IPR is a hybrid graphical representation in which object-level behavior is represented as a Petri net and method-level behavior is represented as function nodes and their data dependency relationships. Correctness preserving transformations rules from the PROOF/L program to the IPR and from the IPR to a target language have also been developed. The two-level transformation via the IPR facilitates the separation of the semantic issues from the performance oriented issues.
- A two-level allocation approach has been developed. At the object level, we model the PROOF program as a directed graph by analyzing the body of each object and then cluster the objects so that each cluster can be analyzed separately at the method level analyzer. At the method-level, various code transformation techniques for various types of parallelism have been developed and compared with other existing approaches using examples. In order to analyze the gain sizes, we develop the notion of gain with which the possible modification to the reduction of the completion time of the program can be determined.

CHAPTER 1 PROOF COMPUTATION MODEL

The object-oriented paradigm is motivated providing for the development of schemes for parallel processing systems [33, 34, 35, 36]. Such a paradigm is based on the principle of information hiding. The main mechanisms associated with the object-oriented paradigm are data encapsulation and inheritance. These mechanisms offer many advantages, such as comprehensibility, modularity, maintainability, reusability and extensibility. In addition, the object-oriented paradigm can naturally reflect the structure of the problem space. In traditional object-oriented computation models, there is always one active object at a time. The active object can send a message to a passive object. Then the passive object becomes active and the sender object waits. After the passive object returns a result back to the sender object and becomes passive, the sender object can resume its computation.

This model is inherently sequential since it was developed for sequential computation. Some attempts have been made to incorporate parallelism in the object-oriented paradigm. In POOL-T [4], parallelism is achieved by allowing execution of the objects without receiving a message. Concurrent Smalltalk [38], Hybrid [35] and languages based on the Actor model [3], such as Act-4 [39] and Scaetta [40], introduce parallelism by adopting asynchronous message passing. The actor model supports delegation instead of inheritance. However, in these languages, only parallelism among the objects can be expressed and their exploitation of fine grain parallelism is difficult or expensive. Our computation model is an approach for exploitation of various levels of granularity of parallelism based on object-oriented and functional paradigms.

In PROOF, a program is represented as a set of the objects which can be executed in parallel. Each object is an instance of a class, and each has its own local data and methods. The methods in PROOF are defined as purely applicative functions. Parallelism at different levels of granularity can be exploited. The major features of PROOF are

- Class, object, and inheritance are supported in PROOF with the restriction that all the methods in the objects are applicative functions. In other words, a functional paradigm is adopted at the level of method definition.
- The goal of the method is structured to support the synchronization between the concurrent objects. Furthermore, it will be shown that the methods along with their guards are calculable.
- Objects are persistent in PROOF. The reception of values by the objects is introduced to modify the objects.

These features are discussed in detail in the following. This chapter is organized as follows. In Sections 3.1–3.3, the features of PROOF are presented with examples. In Section 3.4, the sources of parallelism in PROOF and its semantics are presented.

3.1 Classes and Objects

Class interface and definition

In PROOF, a program consists of a set of the objects. Every object is an instance of a class. A class is a template for a set of the objects having similar behavior, and it is defined as a generic abstract data type. A class is defined by its interface and definition. The class interface describes the types of the methods provided by the class. The class definition consists of the composition of local data and methods, which are purely applicative functions.

```

class interface LimitedBuffer {
  datatype, size
  method get : LimitedBuffer -> datatype -> LimitedBuffer
  method put : LimitedBuffer -> LimitedBuffer * datatype
  method is_empty : LimitedBuffer -> bool
  method length : LimitedBuffer -> int
end class

```

Figure 2.1 The interface of the class `LimitedBuffer`

It is important to point out that restricting the methods of classes application functions does not reduce the expressive power of PROOP, but merely requires that all the effects of the methods be explicitly specified via the parameters of the methods.

Example 4. Class interface and definition of `LimitedBuffer`

We introduce the `LimitedBuffer` example to illustrate the features of PROOP. The class `LimitedBuffer` is a FIFO queue of a limited capacity. The class interface and the definition of `LimitedBuffer` are shown in Figures 2.1, and 2.2. `LimitedBuffer` has two parameters: `datatype` and `size`. `datatype` indicates the type of elements to be stored in the buffer and `size` specifies the capacity of the buffer. The method `is_empty` are self explanatory. In the class definition, the composition class defines the composition of the local data of the class. The composition of `LimitedBuffer` is a Cartesian product. The components can be referred using the dot syntax, such as `b.w` where

Inheritance and genericity

Both inheritance and genericity are supported in PROOP. Classes in PROOP are related by inheritance relations. Inheritance is used to define a subclass as a specialization of a superclass. In a subclass, all the local data and the methods of its superclass are inherited. Additional local data and new methods may be introduced. The inherited methods may also be overridden by a new definition of the method.

```

class definition bounded.Buffer(ctype: IntType, size:
  Int) {
  constructor
  store [Int(ctype)] x = readInt
  method put: Int =>
    expression
    if (appendRight x), then D
  method get: Int =>
    expression
    [ if (isL, then) [x], handle store ]
  method isEmpty: Boolean =>
    expression
    ! count > 0
  method length: Int =>
    expression
    count
end class

```

Functions used here are defined as follows:

```

[appendRight x] x = p :: x
let [p0, p1, ..., pn] = [p0, p1, ..., pn]
let [p0, p1, ..., pn] = r1
size(p) = n + 1
dec(p) = n - 1
f(p0, ..., pi, [pi+1, ..., pn]) = [f(p0, ..., pi, pi+1), ..., f(pi, ..., pn)]

```

Figure 3.3. The definition of the class `Bounded.Buffer` without guard constructs

Example 3. Instantiation of classes.

The class *inheritance* and definition of `Extended.Buffer`, which is a subclass of `Bounded.Buffer` is shown in Figure 3.4. In `Extended.Buffer` a new LIFO method, `pop`, is added, and all the methods in `Bounded.Buffer` are inherited. \square

Generativity is used to define parameters or generic classes. An instantiation of a generic class is obtained by assigning values to the parameters of the generic class. All the local data and the methods of the generic class are inherited. The

```

class interface ExtendedBuffer (stsubtype, sized)
  inherit put, get, is_empty, length from ExtendedBuffer
  method pop: ExtendedBuffer → ExtendedBuffer is stsubtype
end class

class definition ExtendedBuffer (stsubtype, sized)
  superclass: ExtendedBuffer(stsubtype, sized)
  inherit put, get, is_empty, length
  method pop is
    expression
    { if (cast dec) is, then is sized }
end class

```

Figure 3.3: The definition of the class `ExtendedBuffer`

`ExtendedBuffer` in Example 1 is a *generic class*. An instantiation of `ExtendedBuffer` is the following:

```
class My-Buffer instance of ExtendedBuffer (my-type, my-sized)
```

Class `My-Buffer` is an *instantiation* of the generic class `ExtendedBuffer`.

Active and passive objects

A program in FBCOP consists of a set of *objects*. The objects can be classified into the following three categories: *active*, *passive* and *pseudo-active*.

Definition 3.1.1 The *active object* is an object which can initiate a process by invoking a method or methods without any request from the other object.

Definition 3.1.2 The *passive object* is an object which has methods to be invoked by the other objects and cannot initiate them without a request from other objects.

Definition 3.1.3 The *pseudo-active object* is an object which cannot initiate a process but can invoke the methods defined in another object upon request by other objects.

```

class interface ProducerClass {
  method produce() --> ItemType
end class

class interface ConsumerClass {
  method consume() --> ItemType
end class

```

Figure 3.4: Interfaces of the classes `ConsumerClass` and `ProducerClass`

Definition 3.1.1 The *non-passive object* is an object which is not passive, and thus includes active and pseudo-active objects.

Definition 3.1.2 The *non-active object* is an object which is not active, and thus includes pseudo-active and passive objects.

A non-active object acts like a service agency. It waits passively until one of its methods is needed by other objects. The non-active object may in turn provide methods to other objects. An active object is active initially, and it may remain active throughout the execution except for occasional suspensions for the purpose of synchronization with other objects. A body will be attached to each non-passive object. The bodies of the objects are functions that may be recursive and diverge (not terminating).

Example 3: Producer-consumer problem

Interfaces of two classes, `ConsumerClass` and `ProducerClass`, are shown in Figure 3.4. The method `produce()` is to generate an item every time it is called, and the method `consume()` is to consume an item every time it is called. The objects in the producer-consumer problem are shown in Figure 3.5. `Producer` and `Consumer` are active objects. `Producer` will instantaneously generate items and `Consumer` will

Object Buffer	Instance of Buffer Class (qy-1000, qy-1000)
Active Object Body	Producer Instance of Producer Class (qy-1000) /* a body is attached to Producer */
Active Object Body	Consumer Instance of Consumer Class (qy-1000) /* a body is attached to Consumer */

Figure 3.1: The set of the objects for a producer-consumer problem

mutually inconsistent. The definition of the bodies of **Producer** and **Consumer** will be given later. **Buffer** is a passive object, and it becomes active when its methods are invoked by **Producer** or **Consumer**.

3.2. Method Definition

Methods in **PROOF** are purely applicative functions or functional forms, i.e. high order functions. We use a constructor $[n_1, n_2, \dots, n_k]$ to denote a sequence of homogeneous or heterogeneous elements. In the case of homogeneous elements, it denotes a lattice or array whose types are T^* and T^* ($n \in \overbrace{T \times T \times \dots \times T}^k$) respectively. In the case of heterogeneous elements, it denotes a *Cartesian product* whose type is $\prod_{i=1}^n T_i$ ($n \in T_1 \times T_2 \times \dots \times T_k$).

PROOF supports a set of primary functions, such as arithmetic operations, logic operations and list handling operations and functional forms from which other functions and functional forms can be easily constructed. The following are the functional forms currently supported in **PROOF**.

- a) Functional form: n (called *apply to all*)

$$\begin{aligned}
 & n_f[n_1, n_2, \dots, n_k] \\
 & \equiv [f(n_1), f(n_2), \dots, f(n_k)]
 \end{aligned}$$

α has two parameters, a function of type $T_1 \rightarrow T_2$, and a list of homogeneous elements of type T_1 . The function f is applied to each element in the list and yields a list of elements of type T_2 .

b) Functional form: f (called *distributed apply*)

$$\begin{aligned} &f(f_1, f_2, \dots, f_n)(x_1, x_2, \dots, x_n) \\ &= [f_1(x_1), f_2(x_2), \dots, f_n(x_n)] \end{aligned}$$

f has two parameters, a list of functions in which each function f_i is of type $T_1^{(i)} \rightarrow T_2^{(i)}$, and a list of homogeneous elements α in which the i th element is of type $T_1^{(i)}$. Each function in the first list is applied to the corresponding element in the second list. It yields a list in which the i th element is of type $T_2^{(i)}$.

c) Functional form: γ (called *filter*)

$$\begin{aligned} &\gamma(b_1, b_2, \dots, b_n)(x_1, x_2, \dots, x_n) \\ &= \begin{cases} [x_i] & \text{if } n = 1 \text{ and } b_1 = \text{False} \\ [x_i] & \text{if } n = 1 \text{ and } b_1 = \text{True} \\ \gamma(b_1, \dots, b_k)(x_1, \dots, x_k) \cup \gamma(b_{k+1}, \dots, b_n)(x_{k+1}, \dots, x_n) & \text{if } n > 1 \text{ and } 1 \leq k \leq n \end{cases} \end{aligned}$$

First, \cup denotes the concatenation of two lists. γ has two parameters, a list of booleans and a list of any elements. This function yields a subsequence of the second list by selecting elements whose corresponding elements in the first list are *True*.

d) Functional form: f (called *mapped apply*)

$$\begin{aligned} &f(f_1, f_2, \dots, f_n)(x) \\ &= [f_1(x), f_2(x), \dots, f_n(x)] \end{aligned}$$

f has two parameters, a list of fun ctions in which each function f_i is of type $T \rightarrow T_2$, and an element α of type T . Each function in the first list is applied to the element and the list of elements having type $T_1 \times T_2 \times \dots \times T_n$ is yielded as a result of f application.

e) Functional form: η (called *mapped apply*)

$$\begin{aligned} & q(f)(x_1, x_2, \dots, x_n) \\ &= (f(x_1, f(x_2, \dots, f(x_{n-1}, f(x_n, 0)))) - 1 \end{aligned}$$

q has two parameters, a function of which type is $T \rightarrow T$ and a list of elements in which each element is of type T . The function is first applied to the last two elements. Then the result of that application and the third from the end of the list are used as input to another function application. The application is continued until the first element of the list is input to the function. In fact, this functional form can be given a different meaning by allowing parallel application of the function to the elements. That is, we can group the elements into a set of pairs and apply the function f . By repeating this process until a value is returned as a result, we can exploit true parallelism.

The following example illustrates the parallelism due to parallel evaluation of arguments and functions.

Example 4. Partition an array of integers, A which has type int^n , using a pivot element $x \in A$ as shown in Figure 3.4. In other words, we want to rearrange the elements in the array A so that all the elements whose values are less than the value of x precede the elements whose values are greater than or equal to x .

A solution to the problem can be formulated as follows: First, we define the following two functions f and g :

$$\begin{aligned} f &: \text{int} \rightarrow \text{bool} \\ f(x) &= x < x \end{aligned}$$

$$\begin{aligned} g &: \text{int} \rightarrow \text{bool} \\ g(x) &= x \geq x \end{aligned}$$

One solution to the problem is the following non-heap:



Figure 3.4. An array partitioning problem

$$(\gamma \text{ (} \lambda x. f\ x) \ A) \ \& \ (\gamma \text{ (} \lambda x. g\ x) \ A)$$

The functions are curried and we assume left association. This expression is a re-evaluation of two γ functions, which can be evaluated in parallel. Within each of the γ functions, the application of f and g to A can also be evaluated in parallel owing to the functional form α . The execution of this example is illustrated in Figure 3.7.

□

3.3. Synchronisation of Objects

The methods defined on each object may require some *preconditions* under which they can be executed. For example, `get` on `FixedBuffer` can be executed only when the buffer is not empty. When the buffer is empty, the invocation of the method must be suspended. This problem is commonly known as a *synchronisation problem*.

In order to specify the synchronisation constraints in the parallel object-oriented paradigm, many constraints have been proposed. Among them are critical sections [36], message queues [34], behavior abstraction [35] and method-sets [37]. However,




Figure 3.7 A solution for the array partitioning

the use of these constructs interferes with the inheritance mechanism, and the synchronization constraints are not directly inheritable with the methods. In PICOOP, synchronization among the objects is achieved by attaching an optional predicate, called *guard*[30], to each of the methods in a class. Each guard is a predicate. The object which invokes the method is suspended when the attached guard evaluates to *False*, and it is resumed when the guard becomes *True*. The guard attached to the method is defined in a way that it depends only on the status of the local data, and does not on the definition of any other methods. Therefore the inheritance of

```

class definition boundedBuffer(intType, size)
  composition view = list(intType) = view() and
  method put is a
    guard {b.count < size}
    expression
      if (opposite.right a(), list) is
  method get is
    guard {b.count > 0}
    expression
      { if (next, dec) is, last(b.view) }
  method is empty is
    expression
      b.count = 0
  method length
    expression
      b.count
and class

```

Figure 14: The definition of the class `boundedBuffer`

individual methods will not be hampered by the inclusion of the guard. Such guards can be inherited with the methods they are attached to.

Example 5: Inheritance of synchronization constraints

The complete definition of the class `boundedBuffer` with guards is shown in Figure 14. The guard attached to the method `get`, $\{b.count > 0\}$, indicates that `get` can be executed only when the buffer is not empty.

The guard attached to `put`, $\{b.count < size\}$, indicates that `put` can be executed only when the buffer is not full. The complete definition of `ExtendedBuffer`, which is a subclass of `boundedBuffer`, is shown in Figure 15. The `ExtendedBuffer` inherits the methods including the guards of the `boundedBuffer`. The inherited guards work coherently with the new guard defined in `put`. \square

```

class definition ExtendedBuffer(istype t, size)
  superclass: ExtendedBuffer(istype t, size)
  fields: get, get, is empty, length
  method proc k
    guard(s: count > 0)
    expression
      [ if fresh, then 0, length stored ]
end class

```

Figure 3.3 The complete definition of the class `ExtendedBuffer`.

This example has been used to illustrate the difficulty in integrating inheritance and synchronization constraints in some object-oriented parallel models [46, 56]. The above example demonstrates inheritance can be integrated with synchronization constraints without difficulty in *PROOF*. By separating synchronization constraints (guards) from the behavior of methods (expressions), integration of inheritance and parallelism can be achieved with ease. Only modification of the bodies is required.

Since each method is associated with a guard specifying the condition under which the method can be executed, each method can be a natural unit of inheritance.

14. Framework of Objects

A major deficiency of the functional paradigm is its history insensitivity. *PROOF* is made history sensitive by making the objects persistent and allowing the reception of values by the objects, i.e. the assignment of values to the objects. The local data of the objects is persistent. The reception of values by the objects will modify the local data of the objects. A partial function \mathcal{R} , called the reception function, is introduced to denote the reception of a value by an object.

$$\mathcal{R}[\bullet](x)$$

Object <i>Buffer</i>	Instance of <i>Unshared.Buffer</i> (<i>xy-size</i> , <i>xy-max</i>)
Active Object Body	<i>Producer</i> : instance of <i>Producer.Class</i> (<i>xy-size</i>) <i>state</i> : <i>True</i> (<i>R</i> [<i>Buffer</i>] (<i>get Buffer</i> <i>producer</i>)
Active Object Body	<i>Consumer</i> : instance of <i>Consumer.Class</i> (<i>xy-size</i>) <i>state</i> : <i>True</i> (<i>R</i> [<i>Buffer</i>], <i>consumer</i>) (<i>get Buffer</i>)

Figure 3.10: The definition of the objects for producer-consumer problem

R is not a function, but can be treated as such. *R* has two parameters: an object *x*, the recipient, and the expression *e*, to be received by *x*. *x* may contain applications of applicative functions only. This pseudo-function can only appear inside bodies of active objects, and may not be nested. Major differences between modification of the objects through *R* and traditional composition are:

- The evaluation of the expression *e* in *R* can be done in parallel, since *e* contains only applications of purely applicative functions.
- No partial modification to the object *x* is allowed. The local data of an object can only be modified as a whole entity, i.e. its components cannot be modified individually.

The restriction that the recipient functions can only appear in the bodies of the objects implies that all the methods in an object are still applicative, i.e. there are no side-effects.

This restriction effectively preserves the referential transparency at the method level, and retains the parallelism owing to that referential transparency.

Example 8: Bodies of active objects

The complete definition of the objects for a producer-consumer problem is shown in Figure 3.11. In the body of *Producer*, *put* adds an item produced by the method

produces the buffer, and returns a new buffer as its result. The object buffer will receive the value of the new buffer. In the body of `Consistent`, `get` returns the value of a new buffer with an item deleted from the original buffer and the deleted item. The buffer will receive the value of the new buffer and the deleted item will be returned by the `Consistent` as the result of invoking `consistent`. □

Simultaneous access to the objects

An object can participate in more than one function evaluation. There may be several attempts to modify the same objects at the same time. Simultaneous modification of the objects may result in inconsistent and incorrect states. Thus, simultaneous modification to the same object must be controlled so that at any given moment an object will be a recipient of only one of the function evaluations. In order to guarantee simultaneous access to the objects, let's consider the methods defined in the `ArithmeticBuffer` shown in Fig. 2.8. Simultaneous invocation of `length` and `put`, or `length` and `get` are permitted since only one of the methods will attempt to modify the object. On the other hand, simultaneous invocation of `put` and `get` are prohibited since both the methods will attempt to modify the object. Thus, they must be controlled.

At any moment, the status of any object involved in an expression falls into one of the following three categories:

- read only** : The expression only needs to read the value of the object.
- will modify** : The expression will modify the object, but the modification does not occur at this moment.
- modifying** : The expression is currently modifying the object.

In order to ensure the consistency and the correctness of the objects, a multi-mode locking mechanism is adopted. There are three different types of locks, R-Lock,

Table 3.1. Kiehl's mode locking mechanism

	R-Lock	W-Lock	M-Mode
R-Lock	compatible	compatible	incompatible
W-Lock	compatible	incompatible	incompatible
M-Mode	incompatible	incompatible	incompatible

R-Lock and W-Lock, that are associated with the three states of the object, read-only, write-only and modifying, respectively. Before evaluating an expression involving an object x , a proper lock for x must be obtained. A lock is granted only when it is compatible with other locks granted for the same object, according to the compatibility chart in Table 3.1.

3.3. Control Functions

The functional forms used to express parallelism are based on data dependency. In other words, the applications of functions can be ordered according to the input and output relations. However, we may need functions whose relations are not directly based on input/output relations. For instance, consider a dining philosophers problem. In this problem each philosopher either thinks or eats spaghetti. In order for a philosopher to eat spaghetti, he first has to acquire two chopsticks. A philosopher repeats a sequence of activities: think, acquire chopsticks, eat spaghetti and release chopsticks. If we consider each activity to be a function, each can be designed as a method in PROLOG. In this case, there is no significant data flow between think and acquire chopsticks. As the programmer knows from the requirements of the problem that every philosopher will first think and then request chopsticks by invoking acquire chopsticks, he has to express such sequential ordering of activities. However, the function think may not have significant output generated and used as input to the function acquire chopsticks. Thus, if we do not have a sequential control function,

we need to create data flow relationally between these two functions. In fact, we can use a kind of flag as an output parameter of a preceding function to indicate the completion of the preceding function, such as *done*. The same flag is used as an input parameter to trigger the following function *enqueue* objects. The use of such flags not only makes programming very complicated, but also reduces readability and maintainability of the program.

To alleviate such problems, we introduce the following functions, called control functions. These are applications which require explicit control structure other than data flow, we need additional constructs to express high level sequential or parallel control structures. γ is a sequential control function defined as follows:

a) Control function: γ (called sequential)

$$\begin{aligned} & \gamma(a_1, a_2, \dots, a_n) \\ & \equiv a_1; a_2; \dots; a_n \end{aligned}$$

in which γ means sequential execution. It is not necessary that there are any data dependency relations among a_i .

\parallel is a parallel construct defined as follows:

b) Control function: \parallel (called parallel)

$$\begin{aligned} & \parallel(a_1, a_2, \dots, a_n) \\ & \equiv a_1 \parallel a_2 \parallel \dots \parallel a_n \end{aligned}$$

in which \parallel means parallel execution. These control functions are used to specify sequential control flow or explicit parallel control flow among functions in the bodies of the objects.

In addition to these explicit control functions, we also define *while* and *if* control functions as follows:

a) Control function: *if* (called *conditional select*)

$$\begin{aligned} \text{if } [p \ f \ x] \\ \rightarrow \begin{cases} f(x) & \text{if } p(x) \text{ is true} \\ x & \text{if } p(x) \text{ is false} \end{cases} \end{aligned}$$

in which p , f and x are of type $T \rightarrow \text{boolean}$, $T \rightarrow T_1$, and $T \rightarrow T_2$, respectively, and x is of type T

b) Control function: *while* (called *conditional loop*)

$$\begin{aligned} \text{while } [p \ f \ x] \\ \rightarrow \begin{cases} \text{while}[p \ f](f(x)) & \text{if } p(x) \text{ is true} \\ x & \text{if } p(x) \text{ is false} \end{cases} \end{aligned}$$

in which p and f are of type $T \rightarrow \text{boolean}$, $T \rightarrow T_1$, respectively, and x is of type T . The control functions, *if* and *while* are applicative functions without side-effects. They can be used in the method definition as well as in the body.

2.4. Parallelism in PROOF

In this section, we discuss the parallelism offered by PROOF and its semantics.

2.4.1. Sources of Parallelism

PROOF offers parallelism primarily in two different levels: the object level and the method level. The object level parallelism is achieved by allowing more than one object to be active at a time, and the method level parallelism can be achieved by defining methods as applicative functions.

Each of the active objects can invoke a method, thus there can be a number of methods being executed simultaneously. The potential for the simultaneous participation of an object in the evaluation of different functions also implies another

sources of parallelism in PROOF. Method level parallelism is obtained from the following two sources: parallel evaluation of arguments of functions and parallel evaluation of functions or evaluation of a function. The former is achieved because of referential transparency of applicative functions. The latter is achieved by functional forms, such as α and β . In the following, we summarize the sources of parallelism in the computation model. Because all the functions in PROOF are applicatives, parallelism is obtained from the following two sources:

- P1** Parallel evaluation of arguments of functions. It is made possible because of the referential transparency of applicative functions.
- P2** Parallel evaluation of a number of functions or evaluations of a function. It is made possible by functional forms such as α and β .

The active objects introduce another source of parallelism in PROOF.

- P3** There can be a number of the objects that are active simultaneously throughout the execution.
- P4** An object can be involved in two or more function evaluations simultaneously.

3.1.1. The Parallel Semantics of PROOF

In this section, we will define the semantics of PROOF in terms of its parallel behavior. We give the semantics to PROOF at three different levels: one is at the method definition level (applicative functions), another is at the object definition level and the other is the interface level between the two levels.

Parallel evaluation of arguments and parallel evaluation of a number of functions are related to the issue of the semantics of structured programming languages. At the method definition level, an important concern is how to select the expression to be

executed next. There are three different semantics for programs written with functional programming languages, according to how an expression to be executed next is chosen: *strict*, *lazy* and *lexical*. *Strict* semantics requires the following requirement during evaluation of the expression: complete evaluation of all arguments before a function call and complete evaluation of all components before a structure is built. This requirement implies that strict semantics may require useless computations that is not necessary to be executed. This has two disadvantages: (1) the consumption of resources and (2) a failure to produce a solution. It also requires additional effort to implement in distributed memory parallel computers, since some of the unnecessary results may complicate their handling and increase overhead in communication costs.

In contrast to such strict semantics, *lexical* semantics does not impose any requirement mentioned above. Both *lazy* and *lexical* semantics are described in the sense that any of the above requirements for strict semantics is not required. Instead, *lazy* semantics imposes the following requirement: a computation cannot be executed unless it contributes to the final answer. This means that function arguments, structure components and some of conditionals are not evaluated unless they are needed. Furthermore, termination of the computation and obtaining the final answer occur at the same time, since there will be no further reductions required when a normal form is reached. As a consequence, unlike the case of strict semantics, *lazy* semantics generates minimal computation to obtain a final answer if it exists. However, the *lazy* semantics may require a substantial amount of work to determine the next computation to be executed.

Lexical semantics [7] imposes the following requirement: complete evaluation of the predicate before any use of a conditional is considered for execution. Thus, unnecessary computations are avoided only when they are guarded by appropriate conditions. For instance, consider the control function $if\ p\ /\ p\ /\ q$. Under lexical

semantics, first $p(x)$ will be evaluated. Then, based on the result of the evaluation, either $f(x)$ or $g(x)$ will be chosen to execute. When the expression is executed based on strict semantics, the three subexpressions, $p(x)$, $f(x)$ and $g(x)$, can be executed in parallel, since lowest semantics can be considered as a compromise between strict semantics and lazy semantics. Lowest semantics can eliminate substantial effort in determining the next computation to execute. It can also initiate many computations, and thus increase the opportunities for parallel execution. In PROOP we adopt the lowest semantics due to the advantages for parallel processing. In the following, we give the semantics of PROOP in its parallel aspects. The key element in the definition of the semantics is the parallelizer denoted *Para*, which describes how a program in PROOP can be executed in parallel. The parallelizer *Para* is defined using four evaluation schemes presented in the following order:

Scheme A: evaluation of application functions

Scheme B: evaluation of the pseudo-function *R*

Scheme C: evaluation of control functions

Scheme D: evaluation of a set of the objects which compose a program in

PROOP

It is shown that the parallel execution of a program according to the parallelizer *Para* is equivalent to a sequential execution of the same program.

Scheme A. Let e be an expression consisting of applications of application functions only. *Para*(e) is defined as follows:

Case 1: e is a constant

$$\text{Para}(e) = e$$

Case 2: e is a function $f(r_1, r_2, \dots, r_n)$ except the function g' and while

$$\text{Para}(f(r_1, r_2, \dots, r_n)) =$$

```

parallel
   $v_1 \leftarrow \text{Para}(x_1) \parallel v_2 \leftarrow \text{Para}(x_2) \parallel \dots \parallel v_n \leftarrow \text{Para}(x_n)$ 
parallel
 $v \leftarrow f(v_1, v_2, \dots, v_n)$ 
return v

```

Case 3: x is a function of $[p \mid f \mid g]x$

```

Para( $x$  of  $[p \mid f \mid g]x$ ) is
  If Para( $p(x)$ ) then
     $v \leftarrow \text{Para}(f(x))$ 
  else
     $v \leftarrow \text{Para}(g(x))$ 
return v

```

Case 4: x is a function while $[p \mid f]x$

```

Para(while  $[p \mid f]x$ ) is
  If Para( $p(x)$ ) then
    while  $[p \mid f]x$ 
  else
    return x

```

First, `parallel(x_1, \dots, x_n)` and `parallel` is used to indicate the parallel execution of elements separated by `|`. The Case 3 shows that FSDOF supports the nested execution by selecting only one arm for execution after evaluating a conditional. If we treat the g -related function as a general function and apply it to the execution in the case 3, the strict semantics can be obtained. We assume that `Para(x_n)` terminates.¹ However, it may not terminate, or may not terminate. In such a case, we can use a kind of fair scheduling strategy. That is, we can prevent such non-terminating process from

¹We can easily extend this to non-strict semantics by evaluating f in parallel with `Para(x_n)`.

reducing the system resources by limiting number k for each process. With this position, limited resources can produce the same answers as busy resources with only a bounded amount of extra computation.

Variations of scheme A_1 can be derived for the functional forms α , β , γ , δ and η according to their definitions.

Scheme A_{α} . Let e be an expression $\alpha[f_1, f_2, \dots, f_k]$

```

Param( $\alpha$ ,  $f_1, f_2, \dots, f_k$ )  $\equiv$ 
  parbegin
     $v_1 \leftarrow \text{Param}(f_1)(v_1 \leftarrow \text{Param}(f_2)(v_2 \leftarrow \dots$ 
       $\leftarrow \text{Param}(f_k)(v_k) \right)$ 
  parend
  return  $[v_1, v_2, \dots, v_k]$ 

```

Scheme A_{β} . Let e be an expression $\beta[f_1, f_2, \dots, f_k](v_1, v_2, \dots, v_k)$

```

Param( $\beta$ ,  $f_1, f_2, \dots, f_k$ )( $v_1, v_2, \dots, v_k$ )  $\equiv$ 
  parbegin
     $v_1 \leftarrow \text{Param}(f_1)(v_1) \parallel v_2 \leftarrow \text{Param}(f_2)(v_2) \parallel$ 
       $\parallel v_k \leftarrow \text{Param}(f_k)(v_k)$ 
  parend
  return  $[v_1, v_2, \dots, v_k]$ 

```

Scheme A_{γ} is unified since it can be considered as a special form of A_{β} .

Scheme A_{δ} . Let e be an expression $\delta[f_1, f_2, \dots, f_k](x)$

```

Param( $\delta$ ,  $f_1, f_2, \dots, f_k$ )( $x$ )  $\equiv$ 
  parbegin
     $v_1 \leftarrow \text{Param}(f_1)(x) \parallel v_2 \leftarrow \text{Param}(f_2)(x) \parallel$ 

```

```

    
$$[v_n := \text{Param}_L(\bar{u})]$$

  parallel
  return  $[v_0, v_1, \dots, v_n]$ 

```

Solution 2.4. Let v be an expression $\eta[f](x_1, x_2, \dots, x_n)$

```

  
$$\text{Param}_L(\eta[f](x_1, x_2, \dots, x_n)) \equiv$$

  
$$v := \text{Param}_L(f)(x_1, \text{Param}_L(f)(x_2, \dots, \text{Param}_L(f)(x_{n-1}, \text{Param}_L(f)(x_n, \bar{u}) \dots))$$

  return  $v$ 

```

In the following, we prove the statements to be correct. The correctness of statements $A_0, A_{n_0}, A_{n_1}, A_{n_2}, A_1$ and A_2 is proved in the following lemmas and lemmas.

Theorem 2.4.1 The evaluation of application functions in parallel described in scheme 4 is equivalent to a sequential evaluation of the same functions

Proof of Theorem 2.4.1 Since v consists of applications of application functions only, there is no side-effect. The evaluation order of the components/subexpressions of a function does not affect the result of the function. (Charles Rector theorem [34]) Therefore, it implies that the parallel evaluation of all the subexpressions will yield the same result as the sequential evaluation of the subexpressions in any order.

This theorem shows that the evaluation of application functions in parallel is equivalent to a sequential evaluation of the same program.

Lemma 2.4.1 The evaluation of application functions in parallel described in scheme $A_{n_0}, A_{n_1}, A_{n_2}, A_1$ and A_2 is equivalent to a sequential evaluation of the same functions

Proof of Lemma 2.4.1 Since all the functional forms belong to the application functions, by Theorem 2.4.1 it is proved

PROOF exploits coarse grain parallelism by allowing more than one object to be active at a time. Unlike the evaluation of application functions, parallel execution of the objects involves parallel modifications, which must be controlled in such a manner that only the results equivalent to a series of sequential modifications of the objects are permissible. The locking mechanism used in the concurrency control of data base systems is adopted here to ensure the serializability [33] of parallel modifications of the objects. R-Lock is used to indicate that an expression only needs to read an object (read-only). W-Lock is used to indicate that an expression will modify an object, but not now (will-modify). R-Lock indicates that an expression is currently modifying an object (modifying). The compatibility between these locking modes are shown in Table 3.1.

The modification of the objects is achieved by the parallel-function \overline{R} . The execution of $\overline{R} [x] \mid \overline{R}(x)$ is described in scheme B. First, the following notation is introduced. Let x be an expression:

$$\text{obj}(x) = \begin{cases} 0 & \text{if } x \text{ is a constant;} \\ x & \text{if } x \text{ is an object name;} \\ \{x_{i_1}, x_{i_2}, \dots, x_{i_n}\} & \text{if } x = R(x_{i_1}, x_{i_2}, \dots, x_{i_n}). \end{cases}$$

Scheme B. Let x be an object name and x be an expression. Para ($\overline{R} [x] \mid \overline{R}(x)$) is defined as follows:

Case 1: x is a constant.

```
Para ( $\overline{R} [x] \mid \overline{R}(x)$ ) is
  B1.1: W-Lock( $\{x\}$ )
  B1.2:  $x \leftarrow x$ 
  B1.3: W-UnLock( $\{x\}$ )
```

Case 2: x is an expression.

```
Para ( $\overline{R} [x] \mid \overline{R}(x)$ ) is
  B2.1: repeat
```



```

B2.2  W-Lock( $\{x\}$ )
B2.3  W-Lock( $obj[x] \leftarrow \{x\}$ )
B2.4  evaluate guards associated with the methods in  $x$ 
B2.5  If one or more guards are False
B2.6      W-Lock( $obj[x] \cup \{x\}$ )
B2.7  until all guards are True
B2.8   $t \leftarrow Perm(x)$ 
B2.9  W-Lock( $\{x\}$ )
B2.10  $obj[x] \leftarrow obj[x] \cup \{x\}$ 
B2.11  $x \leftarrow t$ 
B2.12 W-Lock( $\{x\}$ )

```

In case 1, we only need to request an W-Lock for the object x before the modification, and unlock it after the modification. In case 2, an W-Lock is requested for each object involved in the expression e except the object x , for which a W-Lock is requested. Then all the guards associated with the methods in x are evaluated. If not all of the guards evaluate to True, we unlock all the objects and repeat the evaluation until all the guards are True. Then, after evaluating the expression e , we request an W-Lock for an object x and modify the object x . Finally we unlock the object x .

Theorem 3.3.2 *Schema B ensures that the parallel modification of the objects is equivalent to a sequential modification of the objects, i. e. the parallel modification of the objects is serializable.*

Proof of Theorem 3.3.2 In [16], a sufficient condition for serializability is given for the two-phase locking protocol. The two-phase locking protocol requires that the locking process consists of two phases, growing and shrinking phases. In the growing phase, locks are obtained, but cannot be released until we complete this phase. In shrinking

phase, locks are released, and no lock can be obtained, including acquiring of locks. In other words, no lock can be obtained after we start to release locks. Therefore we only need to show that scheme B satisfies the two-phase locking protocol.

Case 1: When x is a constant, only one lock is involved. Thus, it is obvious that it satisfies the two-phase locking protocol.

Case 2: When x is an expression in which one or more objects are involved, the growing phase consists of $RL1$ to $RL4$, and the shrinking phase consists of $RL10$ to $RL13$. In fact, the request statement in the growing phase may request unlocking all of the objects. However, the two-phase locking protocol is not violated, since once the objects are unlocked, we will not enter the modifying region ($RL10$ to $RL13$). When we eventually enter the modifying region, we will not request any lock after we start to release locks. Thus, it satisfies the two-phase locking protocol. Therefore, the parallel modification of the objects in PRCOP is equivalent to the sequential modification of the objects.

We give the notation to the explicit control functions as the following.

Scheme C: Let x be an expression (x_1, x_2, \dots, x_n)

$Parallel_([x_1, x_2, \dots, x_n]) =$

begin

$r_1 \leftarrow Parallel(x_1)$

return r_1

$r_2 \leftarrow Parallel(x_2)$

return r_2

end

$r_n \leftarrow Parallel(x_n)$

```

return  $v_1$ 
endend

```

Proc. seqlength(n_1, \dots, n_k) seqend is used to evaluate the sequential execution of elements expanded by a comma.

Scheme C_{seq} : Let e be an expression $f(x_1, x_2, \dots, x_k)$

```

Proc( $f(x_1, x_2, \dots, x_k)$ ) =
  parbegin
    seqbegin
       $v_1 \leftarrow \text{Proc}(x_1)$ 
      return  $v_1$ 
    seqend
  seqbegin
     $v_2 \leftarrow \text{Proc}(x_2)$ 
    return  $v_2$ 
  seqend
  {
    seqbegin
       $v_k \leftarrow \text{Proc}(x_k)$ 
      return  $v_k$ 
    seqend
  }
  seqend
endend

```

Note the difference between Scheme A_p and Scheme C_{A_p} . In the latter, it is not necessary to wait for all the results. However, in the former, the sequential requires collection of the results before proceeding its computation.

Theorem 3.4.2 The evaluation of the implicit control functions described in subroutines G_1 and G_2 is equivalent to a sequential evaluation of the same functions.

Proof of Theorem 3.4.2 In the case of the sequential control function $\{ \}$, the proof is trivial. In the case of the parallel control function $\{ \}$, let a body of an object α , $\text{body}(\alpha)$ be an expression $\{f(x_1, r_1, \dots, r_n)\}$. When the $\text{body}(\alpha)$ includes only application functions, by Theorem 3.3.1 the parallel evaluation is equivalent to a sequential evaluation. When the $\text{body}(\alpha)$ includes the parallel-function or the parallel-functions, the parallel evaluation is equivalent to a sequential evaluation. Therefore, the evaluation of the control functions $\{ \}$ and $\{ \}$ in parallel is the equivalent of a sequential evaluation of these functions.

Note that Theorem 3.4.1 does not guarantee that the parallel evaluation of the body always gives the same result. Theorem 3.4.2 assures that it is always possible to find a sequential evaluation which is equivalent to a parallel evaluation of the body. For instance, when the body involves more than one modification statement to the same object, (i.e. parallel-function \mathbb{M} whose recipient is the same object), the result varies depending on the execution order of the parallel-function statements. In order to ensure that the parallel evaluation of the body always yields the same result, we can make a restriction on the definition of the body.

Definition 3.4.1 A body is called *safe* if and only if there is at most one modification statement for each object involved in the body.

If a body of an object is safe, all the expressions in that body can be executed in parallel. Thus, parallel evaluation guarantees the same result all the time.

Lemma 3.4.1 The parallel evaluation of a safe body ensures a unique result which is equivalent to a sequential evaluation of that body.

Scheme D describes the execution of a set of the objects, that compose a program in PROOF.

Scheme D Let O be a set of the objects which compose a program, and $\bar{O} = (p_1, p_2, \dots, p_n)$ be a subset of the objects that are non-parallel. We use $\text{body}(p_i)$ to denote the body of a non-parallel object p_i . Note that $\text{body}(p_i)$ is a function which may contain the pseudo-function \bar{O} . The parallel execution of O is defined as follows:

$\text{Para}(O) =$

parallel

$\text{Para}(\text{body}(p_1)) \parallel \text{Para}(\text{body}(p_2)) \parallel \dots \parallel \text{Para}(\text{body}(p_n))$

parallel

Theorem 3.4.4 Scheme D ensures that the parallel execution of a program in PROOF is equivalent to the sequential execution of the program.

Proof of Theorem 3.4.4 According to Theorem 3.4.3, any evaluation of the application function in parallel is equivalent to the sequential evaluation of the same function. According to Theorem 3.4.2, any evaluation of the pseudo-function \bar{O} in parallel is available. According to Theorem 3.4.2, any evaluation of a body is equivalent to a sequential evaluation of the body. Therefore, the parallel evaluation involved in the set of the bodies is also available.

CHAPTER 4 INTERMEDIATE PROGRAM REPRESENTATION FOR PROOF/L

There are several intermediate languages for functional language implementations, such as IF1 [12], P-TAC [6], and Lenn [10]. IF1, which was developed for the implementation of the functional language SISAL, is a hierarchical graph language that describes the dataflow graphs produced from SISAL functions. P-TAC is an intermediate language designed to capture the sharing of computation so that optimization of ID programs can be seamlessly understood and classified during the compilation of ID programs. Lenn, which is an intermediate language for specifying computation in terms of graph rewriting, is not aimed at any specific functional language, and it can serve as an intermediate form between functional and machine languages. In general, these intermediate languages serve as a basis for code optimization analysis.

The Intermediate Program Representation(IPR) language is designed to represent the parallelism in the PROOF/L program and analyze it for the efficient exploitation on various parallel computers. The programming language PROOF/L is a parallel object-oriented functional language based on the computation model PROOF. Once a program is written in PROOF/L, the program is transformed to a target code via a two-level translation. To generate a target code for a specific parallel processing system, a PROOF/L code is first translated into an IPR, then the IPR is translated to a target code. The front-end translation makes all the parallelism in the program explicit. Because this step is machine-independent, it needs to be done only once for each PROOF/L program. The back-end translation is machine dependent and the analysis to obtain efficient target codes can be incorporated at this stage. The

two-level transformation approach is shown in Figure 4.1. This chapter is organized as follows: In Section 4.2, the general characteristics of the IPR is presented. In Section 4.3, the formal definition of the IPR is given using the Petri net. In Section 4.4, the transformation rules from the PROSE/L program to the IPR, and from IPR to the target program are presented.




Figure 4.1 Two-level transformation

3.1 Characteristics of the IPR

The IPR consists of two different types of representation: one is a Petri net and the other is a set of function nodes and their relations, which we will introduce in this section. The semantics for the IPR is also given at the two different levels: object level and method level. The object level semantics gives meaning to the object bodies of the PROLOG program. The method level semantics gives meaning to the function nodes used to represent the methods in the PROLOG program. This two level semantics makes it easy to understand the important issues in parallel programs, such as communication, synchronization aspects without considering the unnecessary details of the program. This separation of the semantics also allows the verification of programs at different levels, and thus the complexity of the understanding and the verification of programs can be significantly reduced. Before we present the formal definition of the IPR, we summarize the characteristics of the IPR as follows:

- Graphical representation

A graphical representation can help conceptual understanding of parallel programs as a construct-by-construct basis. A PROLOG program is represented as a hybrid graph on the two different levels: object-level and method level. At the object level, the IPR is a Petri-net, which is a bipartite directed graph. It is used to represent the method invocation structure of the program. At the method level, the IPR is a directed graph in which each node represents a computation and such directed arc represents precedence relationships between two nodes. The nodes can be divided into two types: computation node and non-computation node. The computation node represents a function receiving input value(s) and generating output value(s). These functions are side-effect free, that is, they always produce the same result when the same input values

are given. The non-computation nodes does not represent a specific function, but they are required to start to represent computation in the IFN. A set of edges between nodes specifies the data dependency between functions.

• Data and control dependency representation

The precedence relationships among computations at the method level are based on the data dependency among them. A computation associated with a node can be executed when the input data associated with incoming node(s) is available, and after its execution the result is available for its outgoing node(s). The hierarchical style of PROOFs provides to extract data dependency relationships easily among the computations. The dependency relationships among objects are based on control dependency in the sense that the method is not invoked by the availability of data but invoked by the necessity of that method execution. Thus, we can also regard this control-dependency among objects as a demand-driven approach. On the other hand, the data dependency among the computations in the method level can be regarded as a data-driven approach. The effect of such separation of dependency relations needs to be studied more.

• Partitioning analysis

The IFN can be used to analyze partitioning and allocation of processes to processors. An important factor to measure the execution time of a program is to determine the proper size of the tasks to be assigned as a sequential code to physical processors. To determine the proper grain size, we need to analyze the tradeoff between communication overhead occurring due to parallel execution of processes. In order to use the IFN for this analysis, we need to annotate the graph with information such as execution time and communication time. The

use of the IPE for goal size analysis and clustering will be explained in the next chapter.

4.2 IPE Definition

The IPE is a hybrid graphical program representation, in which the high-level net list structure is specified as a Petri net and in which the functionality is specified as a data dependency graph. The high-level netlist structure corresponds to the bodies of the objects and the low-level functionality corresponds to the method definitions within each object. The high-level structure is called the object-level IPE and the low-level representation of functionality is called the method-level IPE.

4.2.1 Object-level IPE

The IPE in the object-level is represented as a Petri net [35]. A Petri-net represents the static structure of the system and its dynamic behavior. The static structure is represented as a net structure $R = (P, T, F)$, in which P is a set of places, representing the availability of data or context, T is a set of transitions, representing the activities, and F is a flow relation, representing the dependencies between the places and the transitions. To specify the dynamic behavior of the system, the state of the system is represented by a marking on a Petri net, which is a distribution of tokens over places of the net.

We define the Petri net to be used as the object-level IPE for the FSDOF/S program as the following:

Definition 4.2.1 The Petri net PN is defined as a 5-tuple such that

$$PN = (P, T, F, I, R)$$

in which

P is a set of places,

T is a set of transitions,

$F \subseteq (P \times T) \cup (T \times P)$ is a set of arcs (flow relations),

I is a set of initial places, and

D is a set of final places

We also define prestate and postset of places and transitions in the following.

Definition 4.1.1 The prestate of a transition t is a set of places such that

$$st = \{p | p \in P \text{ such that } (p, t) \in F\}$$

Definition 4.1.2 The postset of a transition t is a set of places such that

$$te = \{p | p \in P \text{ such that } (t, p) \in F\}$$

Definition 4.1.3 The preset of a place p is a set of transitions such that

$$sp = \{t | t \in T \text{ such that } (t, p) \in F\}$$

Definition 4.1.4 The postset of a place p is a set of transitions such that

$$pe = \{t | t \in T \text{ such that } (p, t) \in F\}$$

In PN , there is at least one place p such that $sp = \emptyset$. There also is at least one place p such that $pe = \emptyset$. Let num be a function such that $num(p)$ is number of tokens in a place p .

Definition 4.1.5 The initial place is a place p such that $p \in P$ and $num(p) > 0$ and $sp = \emptyset$.

Definition 4.1.6 The final place is a place p such that $p \in P$ and $pe = \emptyset$.

A marking in the Petri net is changed according to the following transition rules

- 1) A transition t is said to be enabled if each input place in st is marked
- 2) A firing of the enabled transition transition t removes a token from each input

place p in π , and adds a token to each output place in π .

When transitions represent computations that are too complex or places that require significant transformations, the transitions or the places can be further refined. Let the refinement of a transition be a sub-net π . The transition can be refined according to the following rules:

- The incoming links and the outgoing links of the transition serve as input and output parameters, respectively.
- There is only one transition that receives all the input parameters.
- There is only one transition that produces all the output parameters.
- All the transitions except these two transitions and places can interact only with the places and the transitions defined within the sub-net π .

The place can be refined according to the same rules as those except that the transition is replaced with the place.

There are two approaches for graphically representing programs: a task interaction graph and a task precedence graph. In the task interaction graph, a program is represented as a undirected graph in which each node represents a task and each edge between two nodes represents a communication relation. Because the edges are undirected, the direction order relation among the tasks cannot be predicted. The known or estimated communication times can be specified by lumping together all of the communication times required during the program execution between processors.

The task interaction graph is suitable for an analysis of task allocation whose goal is to minimize the total execution time and communication time. In the task precedence graph, a program is represented as a collection of tasks and explicit execution dependencies imposed in the form of precedence relationships. This representation is useful for modeling a program whose behavior can be statically known in advance.



Figure 4.2: Procedure relations among functions

In both graphical representations, the concept of the shared data does not exist, and even simple synchronization cannot be specified. For instance, suppose that we have three functions, f_1 , f_2 and *shared_func*, whose relations are shown in Figure 4.2. One interpretation for the relations shown in Figure 4.2 is that the function *shared_func* requires both the results of f_1 and f_2 as input in order to execute *shared_func*. We call this dependency relation as *value flow*. However, there may be situations such that *shared_func* may need only one input from either of the two functions. Suppose that f_1 and f_2 are clients requesting a service in *shared_func*. If *shared_func* can serve only one client at a time, the task procedure graph cannot represent this kind of relations among tasks. We call this relation as *shared flow*.

In the task interaction graph, each relationship could be explicitly represented. However, due to the lack of a procedure relationship, analyzing using procedure relationships cannot be performed on that graph. On the other hand, the Petri net can easily distinguish the differences among these cases, as shown in Figure 4.3. The Petri net is chosen due to its capability to distinguish the thread flow and the value flow. The problem with the Petri net is that recursive behavior cannot be given the



Figure 4.1 The Petri net representation of (a) value flow (b) thread flow

each remainder. Addition of the allocation unit to the Petri net, however, can increase the express power so that the execution of the recursive behavior can be exactly represented.

4.1.2 Method-level IPB

The method-level IPB is a data dependency graph. In the object-level IPB, each method is represented as a transition. In the method-level IPB, functionality corresponding to each transition in the object-level IPB is represented as a set of nodes and their dependency relations. We present function nodes of the method-level IPB and give the semantics using the Petri net. We define the method-level IPB in the following manner:

Definition 4.1.2 The method-level IPB is a directed graph G such that

$$G = (V, E)$$

in which

$v_i \in V$ is a function node, and

$(v_i, v_j) \in E$ is an edge representing data dependency from v_i to v_j .

The method-level IFN is composed of the following nodes:

- A *simple function node* $v_f \in V$ represents a primitive function such as $x_i, x_i + 1, \dots$.

The input and output for v_f can be represented as follows:

$$v_f : I_1, I_2, \dots, I_m \rightarrow O \text{ in which } m \geq 1$$

- A *constant node* $v_{con} \in V$ represents a constant value generator, which produces the specified *const. value(s)* all the time. There is no input to this v_{con} type of node.

$$v_{con} \rightarrow O \text{ which } O \text{ is a constant value.}$$

- An *id node* $v_i \in V$ represents an identity function, which always returns input as output.

$$v_i : I \rightarrow O \text{ in which } I = O$$

- A *copy node* $v_{cp} \in V$ represents a duplicator, which receives an input and produces the appropriate number of copies having the same value as that input.

$$v_{cp} : I \rightarrow O_1, \dots, O_m \text{ in which value of } I = \text{value of } O_i \text{ for } 1 \leq i \leq m$$

- A *main function node* $v_{mf} \in V$ represents a composed function composed of simple and/or main functions.

$$v_{mf} : I_1, I_2, \dots, I_m \rightarrow O, \text{ for } m \geq 1$$

- A *selecter node* $v_s \in V$ represents a conditional construction function. v_s receives input data I_1, I_2, \dots, I_m and control data x and returns an input I_i as an output according to the value of control data x .

$$v_s : I_1, I_2, \dots, I_m, x \rightarrow O$$

- A *distributor* node $n_d \in V$ represents a conditional construction function, which receives input data I and control data c and passes I to one of output port O_i according to the value of c .

$n_d : I, c \rightarrow \{O_1, O_2, \dots, O_m\}$ in which $\{ \dots \}$ means one of the items within $\{ \dots \}$ is chosen.

- A *merge* node $n_m \in V$ represents a nondeterministic selector, which receives an arbitrary number of input data at a time and returns one, which arrives first. If more than one input arrives at the same time, a single input is chosen arbitrarily.

$$n_m : \{I_1, I_2, \dots, I_n\} \rightarrow O$$

- A *split* node $n_s \in V$ represents a decomposition function, which decomposes input data and returns a set of data decomposed.

$$n_s : I \rightarrow O_1, O_2, \dots, O_m$$

- A *concat* node $n_c \in V$ represents a composition function, which composes a set of input data and returns them as one element.

$$n_c : I_1, I_2, \dots, I_m \rightarrow O$$

The selector and distributor nodes were originally introduced in [10], and we generalized them to our method-level IPE. The operational semantics for the nodes in the method-level IPE are given using the Petri net by defining the functionality of the nodes.

The semantics of the basic function nodes, including constant, id, merge function and copy, can be represented as a sample net, as shown in Figure 4-8. The input places indicate the availability of the input value to the function and the output places indicate the availability of the result of the function application. The semantics of



Figure 4-4 The semantics of the function application

the control function nodes, such as *select*, *distribute* and *merge*, and list handling nodes, such as *construct* and *split*, are represented by the Petri net in Figures 4.3-4.7.




Figure 4-5. The operational semantics of the selector node



Figure 4-6. The operational semantics of the distributor node




Figure 4.7: The operational semantics of the list handling nodes. (a) contrast node (b) split node (c) merge node

4.1 Transformation Rules

In this section, we present the transformation rules in the two levels. The PROOF/L program is first transformed into the IPR and then the IPR is transformed to the target code. The former is called *front-end transformation* and the latter *back-end transformation*. The *front-end transformation* is to make all the parallelism explicitly expressed in the IPR. This step is a structure-oriented transformation and machine dependent issues are not involved at all. The *back-end transformation* is performance-oriented and the machine dependent parameters, such as communication overhead, number of links and number of processors, are used to perform various analysis. In the following the transformation rules from the PROOF/L program to the IPR and the transformation rules from the IPR to a target code are given. We also show that the transformation from the PROOF/L program to the IPR preserves the correctness of the original PROOF/L program.

4.1.1 Representation of Interactions among Objects

We transform the behavior of objects into an IPR. The behavior of an object is specified in the body of the object. However, the bodies are defined only for active and pseudo-active objects, not for passive objects. The absence of the bodies for the passive objects causes inconvenience for giving constraints for the passive objects. For detailed information regarding these constraints, refer to [30]. To overcome this drawback, we define the interactions among objects using the following constraints:

- **Sequential execution of methods** When the methods m_1, m_2, \dots, m_n are executed sequentially in the order m_1, m_2, \dots, m_n , its behavior is specified as $SEQ(m_1, m_2, \dots, m_n)$.
- **Concurrent execution of methods** When the methods m_1, m_2, \dots, m_n are executed concurrently, its behavior is specified as $CON(m_1, m_2, \dots, m_n)$.
- **WAIT for method execution** When an object is waiting for the execution of its

method m by another object O is present within its execution, its behavior is specified as $WAIT(m, O)$.

SELECT is a method for execution based on a condition. **SEL** construct behaves like the **CASE** statement in ordinary programming languages. When an object selects one of the methods for execution from among the methods m_1, m_2, \dots, m_n based on a condition, its behavior is specified as $SEL(m_1, m_2, \dots, m_n)$.

ONE-OF the methods for execution from a group of possible methods. **ONE-OF** construct is used in cases where different objects would try to invoke the methods defined in the object O simultaneously. The object O permits only one object to invoke the method at a time. This construct serializes the requests and is typically used to describe the behavior of the shared mutable objects. Note the difference between the **SEL** and the **ONE-OF** construct. Among the set of methods m_1, \dots, m_n defined in an object, when the object permits only one of its methods to be invoked by other objects, the behavior of the object is specified as $ONE-OF(WAIT(m_1, O), \dots, WAIT(m_n, O))$.

4.3.3. Front-end Transformation

We present the front-end transformation rules at two levels: object-level and method-level.

Object-level transformation

Using the **Polis** set defined for the object-level **PP**, we define the transformation rules used in the object-level transformation. To do so, we define the compositional semantics of the operators used in expressing the behavior of the objects with the **Polis** set as follows:



Figure 4.3: Petri net representation of a_1

Let P_1 and P_2 be the Petri net representations corresponding to the behaviors as presented, r_1 and a_2 , respectively, defined as follows

$$P_1 = \{P_1, T_1, P_1, L_1, R_1\}$$

$$P_2 = \{P_2, T_2, P_2, L_2, R_2\}$$

in which P_i, T_i, P_i, L_i and $R_i, i = 1, 2$, are disjoint unless specified otherwise.

We first define a special net, a_1

Definition 4.2.1 a_1 is a net such that

$$P_{a_1} = \{[x], a, a, [x], x\}$$

The net a_1 is a net consisting of only one place without any transition. No transition is possible in a net a_1 . The net a_1 is shown in Figure 4.4.

The semantics of the SEQ construct is given in the following definition

Definition 4.2.2 The transformation rule for $SEQ(b_1, a_2)$ is defined as a net

$$P_{NS} = \{P_1 \cup P_2, T_1 \cup T_2 \cup L_1 \cup P_2 \cup R_2 \cup \{t\} \cup \{t\} \times L_2, L_1, R_2\}$$

The SEQ representation requires an addition of a transition t to sequentially connect the two nets. The sequential composition of the nets, called *assembly*, is shown in Figure 4.5

The transformation rule for the CDF construct is given in the following definition




Figure 4.9: The semantics of the sequential execution

Definition 4.4.3 The transformation rule for $COS(p_1, n_2)$ is defined as a set

$$\begin{aligned}
 \mathcal{TR} = \{ & P_1 \cup P_2 \cup p, T_1 \cup T_2 \cup t, P_1 \cup P_2 \cup \{p, t\} \cup \{i\} \\
 & \cup \{i_1 \cup i_2\}, \{p\}, D_1 \cup D_2 \}
 \end{aligned}$$

The COS representation requires an addition of a place p and a transition t in order to allow the concurrent execution of the two processes represented by each net. The concurrent composition of the two nets is shown in Figure 4.10.

The transformation rule for SEL is given in the following definition.

Definition 4.4.4 The transformation rule for $SEL(p_1, n_2)$ is defined as a set

$$\begin{aligned}
 \mathcal{TR} = \{ & P_1 \cup P_2 \cup p, T_1 \cup T_2 \cup i_1 \cup i_2, P_1 \cup P_2 \cup \{p, i_1\} \cup \\
 & \{p, i_2\} \cup \{i_1\} \cup \{i_2\} \cup \{i_1\} \cup \{i_2\}, D_1 \cup D_2 \}
 \end{aligned}$$

The SEL representation requires an addition of a place p and two transitions i_1 and i_2 so that only one transition can be fired each time. The composition of the two nets for selective execution is given in Figure 4.11.

The transformation rule for ONE-OF is given in the following definition.



Figure 4.12: The semantics of the concurrent execution.

Definition 4.3.1 The transformation rule for $\text{GNE-Diff}(N_1, N_2)$ is defined as a set

$$TR = \{P_1 \cup P_2, T_1 \cup T_2, P_1 \cup P_2, A_1 \cup A_2, B_1 \cup B_2\}$$

in which $|P_1 \cup P_2| = |P_1| + |P_2| - 1$.

Note that there is a common place between the two nets N_1 and N_2 . The GNE-Diff construct is used to represent mutually exclusive access to the shared object. Such access is specified using the predicate function \mathbb{E} in the body of the object. To represent the mutually exclusive access to the shared object, a special place, called the bottleneck place, is associated with the transition representing the method invocation which results the modification of the shared object. The bottleneck place for the object is unique, and thus the number of places after the composition of the two nets is reduced by one. This composition is called *fusion of nets via place*, and is shown in Figure 4.12. The transformation rules presented so far show only the composition of the two nets, but these rules can be easily generalized in the case of the composition of more than two nets.



Figure 4.11: The semantics of the selective invocation

In the computation model *PIRGOS*, a communication is a point-to-point relation in the sense that the communication can only occur between the two objects, the caller and the called. The transformation rule for the object communication is given in the following definition.

Definition 4.3.3 The transformation rule for the communication between two objects is defined as a set

$$\begin{aligned} \mathcal{R} &= \{ P_1 \cup P_2, T_1 \cup T_2, P_1 \cup P_2, A_1 \cup A_2, B_1 \cup B_2 \} \\ &\text{in which } |T_1 \cup T_2| = |T_1| + |T_2| = c, \\ &\quad c \text{ is the number of communication points} \end{aligned}$$

Note in the Definition 4.3.3 that the rule set is an union of two sets, but the two sets T_1 and T_2 are not disjoint. If they are disjoint, there is no method is common in both the sets A_1 and A_2 . It implies that there is no communication between the two objects. The communication between the two objects can only occur when one object invokes a method defined in the other object. The method name is known



Figure 4.12 The creation of the mutually exclusive atom

is relevant to the caller object, and the called object also knows of the existence of the caller object. In the behavior of the caller object, the method name is given with its called object. In the behavior of the called object, the WAIT constraint is explicitly used to specify the possible communication. In other words, the existence of communication between the two objects can be detected by examining the Petri net representation for each body of the two objects. The two nets are composed via the common transition t as shown in Figure 4.12. This composition of nets is called *fusion of nets via transition* and is used only for the representation of communication.

During the transformation, several transitions need to be added to compose the nets. These transitions are not introduced to add functionality to the original nets, but added for the purpose of the composition only. In the following we distinguish



Figure 4.14 The semantics of the communication between two objects

these additional transitions and the transitions corresponding to the functionality of the program.

Definition 4.3.7 The *real transition* is a transition to represent the method invocation.

Definition 4.3.8 The *dummy transition* is a transition used for composition of the nets.

In the transformation rules, the transitions introduced to compose the original nets are all dummy transitions. The firing of the dummy transition does not affect the object state.

Method-level transformation

We present the transformation rules from the PROOF_{FA} method definitions to the method level PN. In the PROOF_{FA} compositional model, each method consists of an optional guard and an expression. Before the expression is executed, the synchronization constraints specified by the guard construct should be satisfied. We present the

semantics of the guard mechanism in the Petri-net. Because each method is represented as a transition in the object level IPR, the semantics for the guard mechanism can be given by relating a transition using the refinement rules given in the previous section. The guard semantics is shown in Figure 4.14 in which the representation of the transition in the object level IPR is shown in the two dotted transition and the boundary of the method is also specified as a dotted rectangle.

The transformation rules for the general function, functional forms and control functions are shown in Figures 4.15 – 4.16. In Figure 4.15, the transformation for a function $f_{2n}(x_1, x_2, \dots, x_n)(x_1, x_2, \dots, x_n)$ is shown in (a), the transformation for the functional form $\lambda f[x_1, x_2, \dots, x_n]$ is shown in (b), the transformation for the functional form $\lambda f[x_1, x_2, \dots, x_n]$ is shown in (c), the transformation for the functional form $\lambda f[x_1, x_2, \dots, x_n]$ is shown in (d), the transformation for the functional form $\lambda f[x_1, x_2, \dots, x_n]$ is shown in (e), the transformation for the functional form $\lambda f[x_1, x_2, \dots, x_n]$ is shown in (f) and finally the transformation for the functional form $\lambda f[x_1, x_2, \dots, x_n]$ is shown in (g). In Figure 4.16, the transformation for the while $[p \rightarrow f][a]$ is shown in (a) and the transformation for the $\text{if}[p \rightarrow f][a]$ is shown in (b).



Figure 4.18: The connection of the method of images with a ground.




Figure 4.15 The transformation rules for function and bounded trees




Figure 8.18 The transformation rules for the control functions

4.1.3. Correctness Preserving Transformations

In this section, we show that the front end transformation is a correctness preserving transformation. In order to show the correctness preserving transformation, we give an interleaving semantics to the PROOF/L program by defining a labeled transition system and a compositional Petri net semantics which can be obtained by applying the front end transformation rules to the PROOF/L program. Then we show the interleaving semantics for the PROOF/L program can be recovered from the Petri net representation of the corresponding IPN form. Note that the Petri net is used to not only give the semantics for the PROOF/L programs but also to give formalism to object level representations in the IPN. In the interleaving semantics, a step of a single process is described as a step of the whole system, changing some global state. Parallel behavior of the system is specified in a sequential interleaved manner. Thus, a computation or program execution is defined as a sequence of such global steps. To give the interleaving semantics for the PROOF/L program we use a labeled transition system [34] as an abstract machine.

Definition 4.1.1 A labeled transition system is a structure A such that

$$A = \langle S, \rightarrow, A_0 \rangle$$

in which

S is a set of states,

$\rightarrow \subseteq S \times Act \times S$ is a transition relation

in which Act is an action to change the state, and

A_0 is an initial state

An element $\langle s, a, s' \rangle \in \rightarrow$ is called a transition (labeled with the action a) from a state s to another state s' by a step a , and will be written as $s \xrightarrow{a} s'$.

In the following, we define a PROOF/L program as a set of processes and then give the semantics to it. Each process corresponds to the body of the non-parallel object. We can associate a set of local states, called S_i , to each object a_i . Then, the state of the PROOF/L program can be represented as a Cartesian product $S_1 \times S_2 \times \dots \times S_n$. The actions to change states are represented by assertions, such as SEL , COR , $ORC-OP$, SEL and $WAIT$. Using the labeled transition system defined in the Definition 4.2.5 and the interpretation of the PROOF/L program given above, the interpreting semantics of the PROOF/L program can be defined as follows.

Definition 4.2.10 An interpreting semantics, A_L , for the PROOF/L program, $PROG$, can be represented by the following labeled transition system

$$A_L(PROG) = (S, \rightarrow, s_0)$$

in which

S is a Cartesian product $S_1 \times S_2 \times \dots \times S_n$,

s_0 is an initial state, and

$\rightarrow \subseteq S \times Act \times S$

in which the transition relations are generated by the following rules

In the rules presented below, $\frac{P}{Q}$ means that if P is true then Q is also true.

Rule 1 Sequential method invocation

$$SELQ(x_1, x_2) \xrightarrow{a_i} x_2, \text{ if } Act \text{ corresponding to } a_i$$

Rule 2 Unconditional selective method invocation

$$\frac{x_1 \xrightarrow{a_i} x_2}{ORC-OP(x_1, x_2) \xrightarrow{a_i} x_2, ORC-OP(x_2, x_1) \xrightarrow{a_i} x_1}$$

Rule 5 Conditional selective method invocation

$$\frac{e_1 \stackrel{\Delta}{=} e_2}{\text{SEL}(e_1, e_3) \stackrel{\Delta}{=} e_2, \text{SEL}(e_1, e_3) \stackrel{\Delta}{=} e_2}$$

Rule 6 Parallel method invocation without communication

$$\frac{e_1 \stackrel{\Delta}{=} e_2}{\text{CON}(e_1, e_3) \stackrel{\Delta}{=} \text{CON}(e_1, e_3), \text{CON}(e_1, e_3) \stackrel{\Delta}{=} \text{CON}(e_1, e_3)}$$

Rule 7 Parallel method invocation with communication

$$\frac{e_1 \stackrel{\Delta}{=} e_1, e_2 \xrightarrow{\text{msg}} e_2}{\text{CON}(e_1, e_3) \xrightarrow{\text{msg}} \text{CON}(e_1, e_3)}$$

The set S of PROOF/L-terms is defined by the following production system:

$$s ::= \text{rel} \mid \text{SEQ}(s, s) \mid \text{CON}(s, s) \mid \text{QSE} = \text{QF}(s, s) \mid s_1 \text{ msg } s_2$$

in which rel is Rel and $s_1 \text{ msg } s_2$ means the communication between two processes via the action labeled msg .

In the following we prove that the first and transformation preserves the semantics of the PROOF/L program by showing that an interleaving semantics of the PROOF/L program is retrievable from its corresponding LPL.

Theorem 4.3.1 The transformation rules given by Definitions 4.3.1 - 4.3.7 preserve the meaning of the bodies of the PROOF/L program

Proof of Theorem 4.3.1 The proof of this theorem consists of the following two steps

1) The *Future* semantics given for the bodies of the PROOF/L program can be

transformed into a labeled transition system.

ii) From the labeled transition system obtained in step i), the interfering semantics for the original PROLOG program can be retrieved.

We call a PROLOG program a PROQ, and let its (PB) representation in the Petri-net be $PN(PROQ) := \langle P, T, P', I, D \rangle$. We first build a labeled transition system LTS from PN as follows:

$LTS(PROQ)$ can be defined as follows:

$$LTS(PROQ) = \{S_0(P), \rightarrow, I\}$$

in which

$S_0(P)$ is a set of non-empty subsets of a set P ,

$\rightarrow \subseteq S_0(P) \times Act \times S_0(P)$ is a transition relation, and

Act is a label corresponding to an action and defined as

$$a \triangleq a', a' \in S_0(P), a' \in S_0(P) \text{ if and only if } \exists B \in T$$

$$\forall p_0, p_1 \text{ such that } p_0 \in I \cap P \text{ and } I \cap p_1 \in P, \text{trans}(p_0) > 0$$

The transition results in $\forall p_0, \text{trans}(p_0)$ is increased by one

We show that from $LTS(PROQ)$ an interfering semantics of the PROQ is retrievable by showing that i) for each useful transition on PN there is a corresponding action in LTS and ii) the synchronization among objects on the PROQ are not excluded.

Next, we show that there are only a finite number of firing transitions introduced during the transformation of any PROLOG program. For each transformation, there are at most two firing transitions added. Thus, for a finite number of the labels in any PROLOG program, we only need to introduce a finite number of firing transitions. The firing and the following useful transitions can be transformed to a new

might introduce without changing the meaning of the original net. Thus, we have proved 1)

Next, from the definition of the transformation rule for communication between two bodies, we know that a transition corresponding to the communication cannot occur unless both the two objects, the caller and the called, are not ready to communicate. Both the caller and called objects can proceed only after the two objects are synchronized for communication (matched connection). We also know by examining the transformation rules given in Definitions 4.3.1 - 4.3.4 that the transformation rules do not violate the matched connection requirement of the original PROOF/L program. Since these transformation rules with the rule for communication are the only rules used for the transformation, the synchronization among the objects cannot be violated. Therefore, we have proved 2)

In conclusion, from 1) and 2), if we give a total ordering to the transitions in $ET\bar{X}$, the total ordering corresponds to an interleaving semantics of the original PROOF/L program.

In the following, we show the method-level transformation also preserves the correctness of the method definition in the PROOF/L program.

Theorem 4.4.1 *The transformation rule for method-level in the ITR preserves the meaning of the methods given in the original definition in the PROOF/L program.*

Proof of Theorem 4.4.1 In the PROOF computation model, each method is defined as an application function without side-effects within the method. The relations among the statements in each method definition are based on data dependency and an applied communication is needed. Thus, it is sufficient to show that the method-level transformation rules do not violate the data dependency relations among functions within the method definition. Since any code in the method-level ITR can be regarded as a

function, by specifying the transformation rule for the general function, it is shown that no violation of data dependency relationship exists. Therefore, we show that the meaning the method in the *PROOF/L* program can be preserved in the representation of the method in the *IPB*.

We can also prove theorem 4.2.2 by building a labeled transition system in which no explicit communication is involved and showing that an interfering scenario can be removed from the semantics of the method level representation of the *IPB*.

Theorem 4.2.2 The transformation rules we introduced for translation of the *PROOF/L* program to its corresponding *IPB* preserve the meaning of the *PROOF/L* program.

Proof of Theorem 4.2.2 From Theorem 4.2.1 and 4.2.1, it is proved

Lemma 4.2.1 The transformation by the transformation rules given in Definitions 4.2.1 - 4.2.2 preserves the correctness of the original *PROOF/L* program.

Proof of Lemma 4.2.1 By the Theorem 4.2.2, the correctness of the *PROOF/L* program is preserved in the *IPB*.

4.2.4. Backend Transformation

Once the *IPB* is generated from the *PROOF/L* program, the various analyses are done in the *IPB* and the modified *IPB* is generated. The analyses includes gross rate determination and clustering, which will be explained in the next chapter. In this section, the translation rules from the *IPB* to a given target language are given.

For the body of each object, the *IPB* is given as a marked Petri net in the front-end transformation. In the back-end transformation, the marked Petri net for each body can be translated to a local scheduler in each processor. A process can be created for each transition in the net and the local scheduler executes the processors in a manner

that the original dependency relationships are not violated. In the case of a network of heterogeneous computers in which the allocation information of the processors should be known prior to run time, the local scheduling information obtained from the object-level IFN needs to be embedded into the target code. Since the net structure can be partitioned, it is not necessary to restrict the boundary of a scheduling unit. That is, an object can be partitioned into a set of parallel processes, and a number of objects can be grouped together to form a larger process. In either case, the net structure can be partitioned or grouped accordingly.

A lock manager needs to be associated to each shared object at this stage. The lock manager controls access to the shared object. In the IFN, the existence of the lock manager is implicit in the sense that the complete semantics of the lock manager is hidden. During the back end transformation, a proper mutex/coarse-grained mechanism can be associated to each shared object.

For each primitive function node in the method-level IFN and control functions, we briefly describe the general transformation rules which are not directly targeted to any specific target language. To introduce general translation rules, we select basic language constructs available in the existing programming languages. These constructs include assignment statement(`var`), while loop statement(`while` ... `do`), conditional statement(`if` ... `then` ... `else`), case statement(`switch` ... `case 1` ... `case n`), function statement(`function`[[`input, output`]]), and process creation statement(`create`[[`process`]]).

In the IFN, every node can be regarded as a function which receives input parameters and which returns a value as a result of function application. The function can be translated to a simple function or a process according to the following rule.

Rule B1: Function translation

Let $func$ be a function. Then $func$ is translated either as a single assignment statement or a *primitive function statement*:

rule B1.1: a *primitive function*

$$Q := func(I_1, I_2, \dots, I_n)$$

in which if $func$ represents *primitive operators*, such as arithmetic

and

logic operators,

$Q :=$ operator I_1 when $n = 1$,

$Q := I_1$ operator I_2 when $n = 2$

rule B1.2: a user defined function

$$create(func(I_1, \dots, I_n, O))$$

The nodes to which the Rule B1 can be applied include *function*, *id*, *constant*, *exp* and *macro function nodes*.

Control function nodes, such as *selector* and *distributor*, are translated according to the following rule:

Rule B2: Control function translation

B2.1: *selector*

when $n = 2$

if $x = true$ then

$$Q = I_1$$

else

$$Q = I_2$$

when $n > 2$

switch x


```

      case 1:  $\bar{O} = f_1$ 
      case 2:  $\bar{O} = f_2$ 
      |
      case m:  $\bar{O} = f_m$ 

  B2.3 decoder
    when m = 2
      if c = true then
         $\bar{O}_1 = f$ 
      else
         $\bar{O}_2 = f$ 
    when m > 2
      switch c
        case 1:  $\bar{O}_1 = f$ 
        case 2:  $\bar{O}_2 = f$ 
        |
        case m:  $\bar{O}_m = f$ 

```

split, *connect* and *merge* nodes are used in the transformation rules for various functional forms as shown in Figure 4.15. In addition to these transformation rules, to utilize the high level programming language constructs, control structures, such as while loop and if then else can also be identified in the IPR.

4.4.4. Example

Using the transformation rules given in this chapter, various programs have been written in PROGF/L and transformed via the two-step translation to the target program written in human C for the simulation of the human transport network

The example programs include factorial, producer-consumer problem, dining philosophers problem, warehouse management systems and air-traffic-control systems. The PROOF/L codes, and the IPL codes for the selected examples are given in the Appendix A2.

CHAPTER 3 ALLOCATION OF PROOFYL PROGRAMS

One of the reasons we execute a program on a parallel processing system is to reduce the time required for completion. It would seem that a simple solution would be to divide all the parallelism in the program and execute it in parallel on the parallel processing system. However, this simple solution does not work in most cases. As the number of parallel processes increases, the communication between processes also increases. Thus, the parallelism needs to be controlled so that the gain of parallel execution cannot be overwhelmed by the communication overheads due to excessive parallelism. We identify this controlling parallelism problem as an *allocation problem*. Various terms used in the literature, such as *allocation*, *partitioning*, *clustering* and *multiprocessor scheduling*, are closely related each other. In general, task allocation used in distributed computing systems, denoting the distribution of tasks into a set of computers. Partitioning or clustering is to group tasks into a set of partitions or clusters, multiprocessor scheduling is to schedule tasks to reduce the completion time. We consider task allocation as a general term which also includes the meaning of clustering, partitioning and multiprocessor scheduling. Whenever further distinction is required, we use proper terms based on the definitions. This chapter is organized as follows. In Section 3.1, we present the existing task allocation approaches, discuss their limitations and briefly outline our approach. In Section 3.2, we present a modeling method for representing the PROOFYL program as a directed graph and a simple bottom up clustering strategy and illustrate it with an example. In Section 3.3, we

illustrate the importance of the *grain size determination*, present strategies for *processor partitioning*, *tree partitioning* and *graph partitioning*, and compare them with the *scheduling approaches*.

3.1 Task Allocation Approaches

The problem of task allocation in a program executes in their *inlined* or *parallel processing systems* has been studied by many researchers, and the existing approaches can be divided into two categories: task allocation without precedence relations and task allocation with precedence relations. Task allocation without precedence relations among tasks can be further divided into three sub-categories: graph theoretic, mathematical programming and heuristic.

In [74], Shao introduced a graph theoretic approach for the task allocation problem in the case of two processors. The program is regarded as a set of tasks, each of which communicates with other tasks. In the program graph, a node represents a task, and an edge between two nodes represents the existence of communication between them. Then the network flow algorithm is applied to this graph. While it is simple to obtain an optimal solution with this method, it can be applied only to a limited number of processors, which severely limits its applicability. An extension of this method for the cases of more than two processors has been proposed as a special case where interprocessor communication patterns are constrained to a tree [75], and a heuristic solution based on this method has also been proposed in [76]. Another method based on a graph has been suggested by Shao in [77]. Shao developed a graph matching algorithm to obtain a task allocation and minimize the completion time of the program when precedence relations among tasks do not exist.

The mathematical programming approach generally involves a 0-1 integer programming technique for mapping program or data flow to processors to minimize

the given cost function with specific constraints [17, 33, 37, 40]. This approach allows constraints to be incorporated into the task allocation model, but is limited by the amount of the time and space required because the time and the space-complexities grow as exponential functions.

While these approaches require expansion of time and space to obtain optimal solutions, heuristic approaches can obtain solutions efficiently by searching the optimality of the solution. A popular heuristic approach is the clustering method [33, 37, 43, 46]. In [36], the proposed method consists of two phases. First, a task clustering algorithm is applied to minimize interprocessor communication. Second, to balance the load, redistributed and overlaid processes are identified and adjusted. In [37, 46], a pair of tasks are searched for clustering such that they clustering eliminates the greatest possible interprocessor communication cost. This process continues until all the pairs are clustered. In [46], two heuristic algorithms have been presented. One is an iterative assignment improvement algorithm. Any initial allocation of tasks is transformed by rearranging tasks to obtain a better allocation. The second algorithm is a clustering method.

In the above approaches we have discussed above, the precedence relations among tasks are ignored. When there are precedence relations among tasks, the goal of task allocation is to reduce the completion time. In general, the task allocation problem for minimizing the completion time is known as the multiprocessor scheduling problem. The multiprocessor scheduling problem has been known as an NP-complete problem except in a few very restricted cases [42, 44]. There have been many heuristics proposed to obtain efficient solutions. Among them list scheduling [3] has been popular because of simplicity and sub-optimality. In list scheduling, the program is represented as a task precedence graph in which each node represents a task and each edge represents the precedence relation between two tasks. Then the list scheduler

assigns a priority to each of the tasks and places tasks in an ordered list according to the priority. When any of the processors is ready to execute, a task with the highest priority is chosen to be executed. Thus, the issue of task scheduling is how to determine the priority of each task.

In [94], the path having the longest length from the entry vertex to the exit vertex, called the critical path, is given the highest priority, and each of the remaining tasks is given a priority based on the number of successors in each task. In [95], WF heuristic has been presented. WF heuristic first gives high priority to tasks that compose the critical path. For the rest of the tasks, composite priorities are calculated based on three criteria: tasks having the longest execution time first, tasks having the largest number of successors first, and tasks having largest immediate tasks first. Although these heuristics are known to be successful, the limitation is that the communication cost between tasks has been completely ignored. In Fom [92], a heuristic approach, called *task clustering*, has been introduced in which not only execution time but also communication cost are considered in scheduling a set of tasks with precedence relations. Erwig[96] introduced a priority algorithm called ETP(Earliest Task First) in which communication times between tasks are considered. ETP adopts the simple heuristic: the earliest schedulable task is first scheduled to an idle processor. However, these two approaches are completely dependent on the given data determined by programmers and thus the performance of these approaches may not be good enough for some applications in which the given data is very small.

In addition, these approaches are not directly applicable to optimize the programs based on the comparative model FBCOP due to the following properties:

- FBCOP offers various granularity levels of parallelism: object-level and method-level.

- The access to the shared object need to be mutually exclusive.

To exploit parallelism in different granularity levels naturally implies a multi-level approach in which a different strategy for each level can be used. The problem with such a multi-level approach is that at the highest level¹ it is difficult to obtain appropriate information for the task allocation. In the PRCOF programs, for instance, accurate execution time at the object level is very difficult to predict because (a) we do not know how much parallelism can be exploited within the methods, and (b) the execution times of the methods are dependent on input data. Another problem is that the access to the shared data may need to be synchronized. However, in most of the existing allocation approaches, the dependency relationship among tasks is based on data dependency, and thus synchronization requirements cannot be specified.

In order to allocate programs based on the computation model PRCOF, we use a two-level allocation approach. At the object level, the parallelism among objects (i.e., coarse-grain parallelism), is exploited. The objective of object-level allocation is to group the objects into a set of clusters so that communication overhead can be reduced while keeping the potential parallelism among the objects. Once the objects are clustered, each cluster has at most one active object. Thus, it is very likely that in each cluster there is one object busy in execution at the moment. We use as input the object invocation relations which can be obtained from the object decomposition phase. At the method level, the parallelism within each method, i.e., fine-grain parallelism, is exploited. At this stage, the proper grain sizes are determined within each method by analyzing the iteration and communication times. Depending on the types of the parallelism, different strategies are required to meet the specific requirement. We present grain size determination strategies for the three patterns of parallelism: page-based parallelism, tree-parallelism and graph parallelism. The grain

¹Target granularity level

our determination strategies can also be used at the object-level if the information about the resources and communication times is explicitly available at that level.

5.2 Object Partitioning

The objective of object partitioning is to partition a set of interacting objects so that the communication cost among processes can be reduced while keeping the potential parallelism among them. The input for our algorithm is the behavior of the objects in the software system specified using the constructs, such as SEQ, COH, QSRGP, SEL and WAIT. The output of our algorithm will be a set of object clusters and their communication dependency relations. The approach consists of the two stages: a modeling stage and a clustering stage.

5.2.1 Modeling

We begin with modeling the software by a directed graph. Each object is represented by one node in the graph, and there is an edge between two nodes if and only if there is a procedure relation between the methods in the two objects. Each edge in the graph has a communication weight denoting the degree of communication between the two nodes. The communication weight associated with an edge represents the communication cost incurred if the two objects represented by the nodes incident to that edge communicate with one another, but they are not allocated on the same processor.

The software system is modeled by a weighted, directed graph $G = (V, E)$. The graph $G = (V, E)$ has a set of nodes V and a set of edges E such that:

- The i th object is represented by a node v_i in V .

- An edge (n_i, n_j) is in E if there is a precedence relation between the two objects such that an execution of a method m_i defined in n_i is followed by an execution of a method m_j defined in n_j .
- A communication weight w_{ij} is associated to every edge (n_i, n_j) .

We assume that the communication weights can be obtained by analyzing the requirement specifications. The factors used to determine the weights can include the frequency of the method invocations and the amount of data transfer required for each invocation.

The PROSPER program is modeled as a directed graph using the following rules. For the simplicity, we use the weight W for each possible invocation when we derive the weights in the following rules.

Rule 1. $n_1 \in \text{CON}\{n_1, n_2, \dots, n_k\}$ describes a case where the objects n_1, n_2, \dots, n_{k-1} and n_k are executed consecutively after being created by the object n_1 . It corresponds to a subgraph $G_k = (V_k, E_k)$ where

$$V_k = \{n_1, n_2, \dots, n_k\},$$

$$E_k = \{(n_i, n_j) \mid 1 \leq j \leq k\}.$$

The communication weights are assigned to the edges as follows.

1. If (n_i, n_j) is new, then $w_{ij} = 1$.
2. If (n_i, n_j) is old, then $w_{ij} = w_{ij} + 1$.

Rule 2. $n_1 \in \text{SEQ}\{n_1, n_2, \dots, n_k\}$ describes a case where n_1 handles n_2, n_3, \dots, n_{k-1} , and n_k in a sequential order. It corresponds to a subgraph $G_k = (V_k, E_k)$ where

$$V_k = \{n_1, n_2, \dots, n_k\},$$

$$E_k = \{(n_1, n_{i+1}) \mid 1 \leq i \leq k-1\}.$$

The communication weights are assigned to the edges as follows.

1. If $\langle x_i, x_{i+1} \rangle$ is *new*, $w_{ij} = 1$
2. If $\langle x_i, x_{i+1} \rangle$ is *old*, $\text{new } w_{ij} = w_{ij} + 1$

Rule 3. $\alpha_1 = \text{ONE-OP}[x_0, x_1, \dots, x_n]$ and $\alpha_2 = \text{SEI}[x_0, x_1, \dots, x_n]$ describe a case where α_1 handles only case α_2 , $2 \leq i \leq n$. Both correspond to a subgraph $G_n = (V_n, E_n)$ where

$$V_n = \{x_0, x_1, \dots, x_n\},$$

$$E_n = \{\langle x_i, x_j \rangle \mid 2 \leq i \leq n\}$$

The corresponding weights in E_n are assigned as follows:

1. If $\langle x_i, x_j \rangle$ is *new*, $w_{ij} = 1/(n-1)$
2. If $\langle x_i, x_j \rangle$ is *old*, $\text{new } w_{ij} = w_{ij} + 1/(n-1)$

Note that instead of the actual method names the corresponding object names are used in these rules. The above three rules can handle simple cases in which only one constraint is used to specify the behavior of the object. The models using these rules are illustrated in the Figure 5.1.

In the following, we present a rule that can handle nested cases.

Rule 4. R_1 is applied when nested cases are used to specify the object behavior. The steps are

- (1) Identify the object behavior by substituting all the nested clauses with new objects
- (2) Select and apply a rule based on the constraint used. For every new object introduced in step (1), do the following steps:
 - (2.1) Apply an appropriate rule and preserve the edge relationships with other objects




Figure 3.1: The modeling of the object behavior

8.8) Assign communication weights as described earlier in this section.

We now have a graph in which each node represents an object (the same as in the initial graph) and to every edge there is now weight w_{ij} representing the semantic proximity score.

3.3.2. Clustering on the Object level

Once the modeling of the software as a directed graph is done, we apply a bottom-up clustering approach to the graph. The main objective of this part is to reduce

the necessary communication overhead by clustering the objects while keeping the potential parallelism among objects.

The input is a weighted, directed graph $G' = (V', E')$. $V' = \{v_1, v_2, \dots, v_p\}$ where v_i is an object for $1 \leq i \leq p$. Every edge $(v_i, v_j) \in E'$ has a weight m_{ij} . The output is a directed graph in which each node represents a cluster and each edge represents the communication requirements between clusters.

The steps we take in clustering the nodes of the graph are as follows.

Algorithm 1.1.1 Object Clustering

input : $G' = (V', E')$

output : a set of clusters

Initially, each object v_i is a cluster $\{v_i\}$.

For each object node v_i having no incoming edge, do the following

Step 1 Set v_i as current object v_c

Step 2 Find a node v_j such that $(v_c, v_j) \in E'$ and m_{cj} is the largest

Step 3 Cluster $\{v_c\}$ and the cluster which v_j belongs to

Step 4 Remove the edges from v_c to the other nodes and from the other nodes to v_c

Step 5 If the node v_j has an outgoing edge to any node then

 Set v_j as the current object v_c

 Go to Step 1

else

 Stop

In Step 1, a node representing an active object is set as a current node to be processed. In Step 2, a node is selected so that the clustering of that node with the current node can reduce the communication overhead most. In Step 3, the node

selected in Step 2 is clustered with the cluster which the current node belongs to. In Step 4 the other nodes connected to the current node are disconnected from the current node. This step guarantees that all possible parallelism in the software is retained by not clustering any two objects having the possibility of parallel execution. These steps are repeated until all the source objects are processed. Because each edge needs to be visited at most once, the time complexity of the clustering approach is $O(e)$ where e is the total number of edges.

The result of the object clustering is a set of clusters, each including at most one active object and the objects reached by that active object. Thus, no parallelism among objects has been lost. Once the clustering of the objects is complete, each cluster is considered as an independent sub program to be analyzed independently of the other subprograms. In each sub-program, each method becomes a subject of analysis. In the following, we illustrate object clustering with an example.

To demonstrate the object partitioning, consider the dining philosophers' problem. One way to implement this problem is to define a philosopher and a chopstick as a class. Each instance of the philosopher class `Philosopher-i` for $1 \leq i \leq 5$, is an active object and each instance of the chopstick class, `Stick-i` for $1 \leq i \leq 5$, is a passive object. The behaviour of the active object `Philosopher-i` can be specified as `{SEQ{lock.acquire-stick(i),acquire-stick(i+1) and S_{i+1} .wait, release-stick(i),release-stick(i+1)}}` in which `'lock'` and `'wait'` are methods in the `Philosopher`, `'acquire-stick(i)'` and `'release-stick(i)'` are methods in `Stick-i` and `'acquire-stick(i+1)'` and `'release-stick(i+1) and S_{i+1} ' are methods in Stick-i+1. The meaning of each method is self explanatory. Since the method invocation defined in the active object is not necessary to specify, the problem can be modeled as shown in Figure 5.2 in which each Philosopher is labeled as P_i and each Stick is labeled as S_i .`



Figure 3.1 An object clustering in the five dining philosopher's problem

When we apply the clustering approach we presented in the previous collection, a possible solution is shown as dotted lines

3.1. Method Partitioning

As we have discussed above, the parallelism must be managed in such a way that the communication overhead can be controlled. The IPR can be used as a task precedence graph by partitioning analysis.

In the method partitioning, the proper grain sizes are determined by recursion and communication-time analysis. In the case of cyclic expressions, such as looping and recursion, we consider only one pass of each cyclic expression. The idea of considering only one pass is to determine the proper grain sizes so that the completion time can be reduced. This one-pass approach can be used to determine the optimal

number of processors. Note that the optimal number of processors remains the same regardless of the number of times the experiment is repeated. We consider this step the case for exploiting small grain parallelism in comparison to the object partitioning step in which large grain parallelism among objects is exploited. In order to perform the grain size analysis based on the tradeoff between parallel execution and communication overhead, we estimate the execution time of each node in the IPB and the communication time between the two adjacent nodes by analyzing the assembly code corresponding to the target code for each IPB construct. In the following sections, we present the existing grain size analysis techniques, and our grain size detection analysis approaches on three different types of parallelism: pipelined parallelism, tree parallelism and graph parallelism.

3.1.1. Grain Size Analysis

In this section, we briefly compare two different strategies in determination of grain size and present our methods on various types of parallelism.

The existing grain size determination strategies can be divided into two categories based on how the grain size is determined: programmer control and automatic determination. In the programmer controlled approach, programmers are fully responsible for determining the grain size, as well as expressing explicitly parallelism. The programmers can use parallel language constructs indicating the tasks to be executed in parallel. When a programmer has specific information about the behavior of a program, the programmer can determine the size of tasks. When that program is ported to a different parallel computer, the sizes of tasks need to be changed to best fit to the new processors. In addition, it may not be easy for the programmers to make decisions on the size of the tasks due to lack of information. This approach

is closely related to the second category of parallel programming approach in which parallel language constructs are used for expressing parallelism and communication.

On the other hand, in the automatic determination approach, grain size is determined automatically. This approach can be further divided into two classes: the compile-time approach and run-time approach. In the automatic determination at compile-time approach, the programmers do not provide any information regarding the granularity. During compilation, heuristics are used to statically determine the size of tasks [35, 38, 70]. One disadvantage is that some information may not be available before the run-time.

In the run-time approach, only simple heuristics can be applied to determine the size due to the costly overhead. For example, as in [34], each recursive function call creates a new task to be assigned to a processor. In this case, since functional programming involves frequent recursive function calls, it is likely that too many small tasks will saturate the system. In general, such run-time approaches ignore the size of tasks under the assumption that there are reasonably many processors available. The automatic determination approach looks more promising since the programmers need not worry about the grain size at all. However, some parallelism may not be detected without help from programmers.

In our approach, we integrate the two strategies by utilizing information available at the compilation time and also allowing programmers' control. The grain size is determined based on the analysis of the creation and communication trees which can be obtained during the compilation phase. The programmer can analyze the IFR and select proper strategies depending on the characteristics of the parallelism.

We make the following assumptions about the underlying MIMD parallel computers:

- The computer system consists of fully connected identical processors having the same processing capability.
- Each processor has a capability of performing program execution and I/O simultaneously.
- The communication cost between two processors depends only on the data size to be transmitted. Currently, we ignore the time required to set up the communication.
- Communication cost between two tasks residing on the same processor is small enough to ignore.

We define the notations we use in this section as the following.

Definition 5.1.1 An execution time for a task α , denoted as $e(\alpha)$, is a time required to complete the computations represented by α without being interrupted.

Definition 5.1.2 A communication time between two nodes α and β , denoted as $c(\alpha, \beta)$ is a time required to transmit data from α to β under the assumption that α and β are assigned to the adjacent processors.

Definition 5.1.3 A completion time for a code α with a number of processors, denoted as $T(n, \alpha)$, is a time required to finish all the computations and communication involved in the code α .

We illustrate the importance of grain size in parallel processing with an example. Suppose that we have a FIBONACCI code, called \mathbf{M}_1 , such as $\alpha[1] \leftarrow 0(\alpha_0)$, $\alpha[2] \leftarrow 1(\alpha_1)$, $\alpha[i] \leftarrow 0(\alpha_2)$, $g(\alpha_3)$. The IPR for \mathbf{M}_1 is shown in Figure 5.2. The graph consists of seven nodes, α , β , γ , δ , ϵ , ζ , η , and six edges, which represent data dependencies




Figure 5.1 An IFE representation for $R1$

relations. This task-type parallelism is a typical form resulting from a divide-and-conquer algorithm.

Assuming that there are enough processors, a simple approach would be to assign four tasks d , e , f and g to four different processors P_1 , P_2 , P_3 and P_4 , respectively. After completing the execution of these four tasks, P_1 and P_2 continue the execution of the nodes b and c , respectively. Then, P_2 executes a to complete the execution. In this case $C(R1, 4)$ can be calculated as follows.

$$C(R1, 4) = 5 + 10 + 2 + 10 + 1 = 28$$

This result is not desirable, since if we only utilize one processor, $C(R1, 1) = \sum_{i=1}^7 c(x_i) = 5 + 4 + 10 + 2 + 1 = 22$. Thus $C(R1, 1) < C(R1, 4)$, and there is no gain in parallel processing because the communication overhead has overshadowed the gain of parallel execution.

3.3.2 Pipeline Parallelism

One of the patterns of parallelism is pipeline parallelism. In exploiting pipelined parallelism, one important consideration is the number of the segments, that is, how to divide the entire process into a set of segments to reduce the completion time. In the following, we show that we can find the optimal size of segments, i.e., the proper grain size.

Definition 3.3.1 A dominant segment is a sub-process in a pipelined process such that the completion time of the entire process is dominated by that sub-process.

Definition 3.3.2 All the sub-processes other than the dominant segment are called subordinate segments.

The dominant segment dictates the completion time of the entire pipelined process. In other words, the dominant segment is always busy once all the segments are filled with data.

Definition 3.3.3 A computation segment is a sub-process whose task is to receive input data, execute the computation with it and return the result.

Definition 3.3.4 A communication segment is a sub-process whose task is to pass data from the preceding computation segment and to the succeeding computation segment.

Definition 3.3.5 A one-pass completion time, ET_{1p} , defined for a segment s_i , is an amount of time required to complete processing of a datum x in s_i .

Note that the dominant segment can be either a computation segment or a communication segment. Computation segment and communication segments have different properties. Computation segments can be divided to reduce the one-pass time.

However, nonoverlapping segments cannot be so relaxed. This distinction implies that in order to find the optimal grain size we need to begin the analysis from the smallest grain size available. In our approach, the IPK generates the needed grain partitions we can get at the program statement level.

We make the following assumptions before we introduce an optimal solution for the determination of the proper grain size.

- The run-time completion time for each segment is fixed and known a priori.
- There are a sufficiently large number of data elements.
- The execution of code and communication can be done simultaneously.
- The communication link can deliver only one result at a time.

Lemma 4.2.1: There is always a dominant segment in a pipelined process.

Proof of Lemma 4.2.1: Under the assumption that each segment requires a fixed amount of time for completing its processing of a datum, the proof is trivial.

The goal of finding proper grain size in a pipelined process is to reduce the run-time completion time of the dominant segment either combining a dominant segment and its neighboring segments or refining a dominant segment. However, since the IPK generates as fine partition as possible, the refinement of the dominant segment would not be considered.

Algorithm 4.2.1: Determine Grains for Pipe-lined Parallelism

input: An IPK representing pipelined parallelism

output: Optimal grain size

Step 1 Sort the requests using the page completion time as a key in descending order.

Step 2 Find a dominant request. Let the dominant request be S_i and its preceding and succeeding requests be S_{i-1} and S_{i+1} , respectively.

Step 3 If S_i is a communication request, then

If $E(S_{i-1}) + E(S_{i+1}) < E(S_i)$ then

Delete S_{i-1} , S_i , and S_{i+1} from the list

Add new S_i into the list such that

$E(S_i) = E(S_{i-1}) + E(S_{i+1})$

Go to Step 1

else

Stop

In the following it is shown that the algorithm 3.3.1 yields an optimal solution.

THEOREM 3.3.1 The algorithm 3.3.1 always finds optimal page sets.

Proof of Theorem 3.3.1. By Lemma 3.3.1, the Step 2 will always find a dominant request. When the dominant request is a computation request, since we cannot reduce the computation any further, the page sets cannot be further reduced. Thus, the dominant request reduces the costs, and so does the completion time of the entire process. Now consider a case in which the dominant request is a communication request. In this case the algorithm tries to partition the dominant request with its neighbor computation requests to reduce the execution page time. Let the sequence of the partitioning obtained by the algorithm 3.3.1 be called P1. Assume that there is another sequence of the partitioning called P2 which can lead to a better solution than

P1 In P2, a communication segment E_i is first partitioned with its neighbors E_{i+1} and E_{i-1} before E_i when

$$E_i(E_i) > E_i(E_{i-1}) \quad (1)$$

Partitioning of E_i with its neighbors E_{i+1} and E_{i-1} means that

$$E_i(E_{i+1}) + E_i(E_{i-1}) < E_i(E_i) \quad (2)$$

From (1) and (2) we know that $E_i(E_{i-1}) + E_i(E_{i+1}) < E_i(E_i)$ (3)

If E_i and E_j are two adjacent communication segments, then in P2 it may not be possible to partition E_i with E_{i-1} and E_{i+1} ($= E_{j-1}$) because $E_i(E_{i+1})$ has increased to $E_i(E_{i+1}) + E_i(E_{j+1})$. Thus, the demand segment E_{j+1} may remain unpartitioned in P2. Therefore, P2 cannot yield an optimal solution. We proved by contradiction that the algorithm SFI can always yield an optimal solution.

Let the number of segments be n . Then, the algorithm SFI can find an optimal group even during compilation in the time-complexity shown in the following theorem.

Theorem 3.2.2 The time complexity of the algorithm SFI is $O(n \log n)$.

Proof of Theorem 3.2.2 Step 1 requires steps steps. By using a priority queue data structure to store the sorted list, we can remove and store any element in steps steps. Thus, each iteration of Step 2 requires at most steps steps for those deletions and one addition to the priority queue. Since there are at most n segments, the entire algorithm can run $O(n \log n)$.

3.2.3 Tree Partitioning

We consider a case in which the task precedence graph is a tree. Before we present our approach, we define the terms we use.

Definition 3.2.3 The root-level sub-tree is a subset of nodes v_0, v_1, \dots, v_n in a tree such that v_0 is a parent node of all the nodes v_1, v_2, \dots, v_n .

In the following, we call the one-level sub-tree simply *sub-tree*. The number of sub-trees in a tree is the same as the number of non-leaf nodes.

Definition 3.2.2 A task precedence tree T_p is a tree in which each node represents a computation and each edge specifies the data dependency relationships among nodes.

Facilities obtained from divide-conquer strategy can lead to the partition of tree patterns and thus be represented by T_p .

Definition 3.2.1 A gain tree T_g of T_p is a weighted tree in which each node, called a gain node, represents a sub-tree in T_p and each edge represents data dependency relations among the nodes. Each gain node has a weight, called *gain*, corresponding to an amount of possible maximum contribution to reducing the completion time when the corresponding sub-tree is clustered.

A gain for a sub-tree consisting of n_1, \dots, n_m is denoted as $GAIN(n_1, \dots, n_m)$. Our gain size determination approach can be considered as a hierarchical clustering or partitioning in that a set of adjacent nodes, i.e., a sub-tree, is considered as a candidate for clustering. The essential part of our gain size determination approach is to estimate the possible contributions which can be made by clustering the adjacent nodes. Our approach consists of two parts:

- 1) build a gain tree from a given input task precedence graph, and
- 2) determine gain size from the gain tree.

The gain tree can be built by analyzing each sub-tree in the task precedence tree using the following procedure.

Procedure Gain Analysis

Input: a sub-tree, consisting of x and its children n_1, \dots, n_m

output: $GAIN(x, n_1, \dots, n_m)$

Step 1 Calculate the total execution time in one processor

$$tt = c(x) + \sum_{i=1}^m c(n_i)$$

Step 2 Find a node n_{i_0} for $1 \leq i \leq m$, such that

$$pt = c(n_{i_0}) + c(n_{i_0}, x) \text{ is the second largest}$$

Step 3 If $pt + c(x) \geq tt$ then

$$GAPN(x, n_1, n_2, \dots, n_m) = pt + c(x) - tt$$

else

$$GAPN(x, n_1, n_2, \dots, n_m) = 0$$

In Step 1, the bottleneck time used to decide whether the computation represented by a sub-tree is or is not executed in parallel is determined by summing all the execution times of the nodes in the sub-tree. In Step 2, the time required to complete the computation represented by the sub-tree when the nodes in that sub-tree are not clustered is calculated. Because at the scheduling stage one of the child nodes can be scheduled to the same processor as the node x and by the static analysis the scheduler can choose a node n_i such that $c(n_i) + c(n_i, x)$ is the largest, the second largest schedule length is calculated to be used as an actual time required to complete the computation represented by x, n_1, \dots, n_m . In Step 3, a gain is calculated by comparing the bottleneck time with the actual completion time calculated in Step 2. If there is a positive gain, then the amount of the gain calculated is assigned to the gain node. Otherwise the gain is set to zero.

The time complexity of the procedure Gain Analysis can be determined by the following analysis. Step 1 requires $O(1)$ time, Step 2 requires $O(m)$ time in which m is a number of child nodes and Step 3 requires $O(1)$ time. Thus, the exact time complexity of the procedure Gain Analysis is $O(m)$.

Now, we build a join tree T_j from a precedence tree T_p using the procedure *Join Analysis* in the following manner:

Algorithm 3.3.1 Build Join Tree

Input: a task precedence tree T_p

output: a join tree T_j

For all sub-trees t_i do

Let t_i consist of a node x and its children m_1, m_2, \dots, m_n

Step 1 Call `Join-Analysis`(x, m_1, m_2, \dots, m_n)

Step 2 Connect the join node to the existing join nodes

In Step 1, *Join Analysis* is called for each sub-tree to determine the possible contribution to the reduction of the completion time when the sub-tree is clustered. In Step 2, the newly created join node for the current sub-tree is connected to the existing join nodes. The construction of the join tree can be done by visiting leaf nodes from the original task precedence tree and associating the join calculated in Step 1 to a parent of the corresponding sub-tree. Step 1 requires $O(m)$ where m is the number of the child nodes. Since Step 1 is executed for each sub-tree and the number of sub-trees is bounded by the number of non leaf nodes in the tree, the algorithm needs to visit each node once. In Step 2, each node also needs to be visited once. Thus, the time-complexity of the algorithm 3.3.1 is $O(n)$ where n is the number of nodes in the tree.

Once the join tree is built, the join size can be determined by selecting clusters hierarchically. Our join size determination is based on the observation that contains more from the nodes close to the root node² propagates to the other nodes. In order

²The root node is a leaf in a node having depth of 0.



Figure 3.4 Sample gain tree examples

To illustrate this, suppose that we have sample gain trees as shown in Figure 3.4 in which $a, b, c \leq 1$, represent a set of gain nodes and a, b, c represent an amount of the gain for each node. In Figure 3.4 (a), the precedence relations are $v_1 \rightarrow v_2$ and $v_1 \rightarrow v_3$. In Figure 3.4 (b), the precedence relations are $v_1 \rightarrow v_2$ and $v_1 \rightarrow v_3$. The goal of the analysis is to select the best two candidates for clustering so as to increase the overall contribution to the reduction of the completion time. The overall contribution is determined by the following equation:

$$\begin{aligned} & \text{overall contribution} \\ &= \max \{ \text{overall contribution from a path between } v_1 \text{ and } v_2, \\ & \quad \text{overall contribution from a path between } v_1 \text{ and } v_3 \} \end{aligned}$$

Thus, when $\max\{b, c\} \leq a$, v_1 and a node having a weight of $\max\{b, c\}$ are chosen to cluster, and when $\max\{b, c\} > a$, v_2 and v_3 are chosen.

Note that the rules presented above can be applied to both gain trees shown in Figure 3.4. It implies that our gain tree analysis technique can be used for the analysis

of both as low (as in as in Figure 3.4 (a)) and not low (as in as Figure 3.4 (b)). The following is an algorithm to determine *good* nodes.

Algorithm 3.3.2 Determine Good Node

Input: A preorder T_p

Output: A clustered leaf preorder tree

Step 1. Initialize all the nodes as 'unclustered'

Step 2. Sort the nodes in T_p using the gain as a primary key in descending order and the depth of each node as a secondary key in ascending order

Step 3. Get the largest node n_i in the sorted list

Step 3.1 If n_i is not a root node and a parent of n_i and

any sibling of n_i are 'clustered' then

go to Step 3.4

Step 3.2 If two or more child nodes of n_i are not 'clustered' then

go to Step 3.4

Step 3.3 Set n_i as 'clustered'

Step 3.4 Delete n_i from the list

Step 3.5 If there is a node with a positive gain then

go to Step 3

else

Stop

The algorithm 3.3.2 determines which sub-tree root is to be clustered by analyzing gains of the possible candidates. Step 1 requires $O(n)$ time to visit each node once. In Step 3, $O(\log n)$ time is required to visit a number of nodes. Step 3 must visit

adjacent nodes at most once for each node. Decrease the number of the adjacent nodes at most in the number of edges in e sets and each edge will be visited at most twice, the overall time complexity in Step 3 is lowered to $O(e)$ in which e is the number of the edges in T . Therefore the time complexity of this algorithm is $O(\max(|\text{edges}|, e))$. Note that in the case of trees the complexity falls to $O(|\text{edges}|)$. Since this algorithm is also used in the case of the graph later, we define the time complexity of this algorithm as above.

In the following, we compare our approach to McCrory's approach [34]. McCrory uses an algorithm of the complexity $O(n^3)$ that decomposes a graph into a set of class that are classified as primitive, linear, or independent. When the class are labeled as independent, the possibility of parallelization exists. However, when the class are labeled as primitive or linear, they are grouped together and executed sequentially. In order to compare our approach to McCrory's, we use the same as explained in [34]. The task precedence tree of the example is shown in Figure 4-4. Every node has a unique number within the circle representing that node, and a weight is attached to the node. The communication cost is assigned to each edge. For instance, a node 3 has a weight 1 and each of the three edges has communication cost 10. Using McCrory's approach, the schedule and its completion time are shown in Figure 4-5.

The problem with McCrory's approach is that it begins to search the candidates for parallel execution from the bottom of the tree. Once the clustering is done at the lower part of the tree, the clustering at the higher levels less likely to occur, since such clustering may have to sacrifice the possibility of parallel execution at the lower level. In addition, the key concept of the graph decomposition approach, when, is determined without using information regarding the execution time and communication cost.



Figure 2.6. A task precedence tree



Figure 5.7. A gain tree for the two-parallel example

Our gain tree determination approach uses the execution and communication times to determine the proper gains at the beginning. By analyzing the contribution of the clustering locally at each node tree, we build the gain tree. Thus, we first select the largest gain in the gain tree as the candidate for the clustering and continue the selection until all the positive gains are processed. The gain tree for this example is shown in Figure 5.7 in which a calculated gain is shown within each gain node. The corresponding nodes in the task precedence tree shown in Figure 5.5 are also associated with each gain node.

From the information obtained in the gain tree, we can determine five groups as follows:

$$G1 = \{ 1, 3, 8, 10, 11, 14, 15 \}$$

$$G2 = \{ 5, 6, 11 \}$$




Figure 5.4 A schedule obtained from our approach

$$C1 = \{ 1, 8, 15 \}$$

$$C2 = \{ 2 \}$$

$$C3 = \{ 4 \}$$

Using these clusters, we show two schedules and their completion times in Figure 5.4. A schedule for five processes is shown (a) in Figure 5.4 and a schedule for five processes is shown (b) in Figure 5.4.

In the following, we present a task allocation strategy which can produce an optimal solution in a restricted case. We first present the algorithm and then show that it can find an optimal solution when the following assumptions are met:

- The task precedence graph is a tree
- There are sufficiently many processors so that the leaf nodes in the task precedence graph can be executed simultaneously, if necessary
- The gain cost analysis has been completed and no more clustering of grains is before automatication has a necessary

This algorithm can be considered as a variant of list scheduling. The priority of the task for allocation to the next processor is decided by the expansion of the task. This algorithm also uses the concept of the gain analysis. The algorithm tries to reduce the completion time of the task node by clustering the critical path with that node.

Algorithm 3.3.4 Task allocation for tree

Input: A task precedence tree

Output: a set of clusters

- Step 1** Identify a node x whose predecessors are all leaf nodes, u_i , for $1 \leq i \leq m$. If there is no such a node x , then stop
- Step 2** If x has only one predecessor, u_i , then replace x , u_i and an edge between x and u_i with a new node x' such that $a(x') = a(x) + a(u_i)$ and go to step 1
- Step 3** For x and u_i , $1 \leq i \leq m$, do the following

Step 3.1 Calculate the total execution time in one processor

$$t_0 = n(p) + \sum_{i=1}^m n(n_i).$$

Step 3.2 Find a node n_{q_1} for $1 \leq q \leq m$, such that

$$n(n_{q_1}) + c(n_{q_1}, n) \text{ is the largest.}$$

Step 3.3 Find a node n_{q_2} for $1 \leq q \leq m$, and $q \neq q_1$, such that

$$n_2 = n(n_{q_2}) + c(n_{q_2}, n) \text{ is the largest.}$$

Step 3.4 If $n_2 > t_0 - n(n)$ then

cluster n and n_{q_2} into a new node p such that

$$p = n(n_{q_2}) \cup n_{q_2} \cup \dots \cup n_{q_1} \text{ and } n(p) = t_0$$

delete n and all n_i from the tree

attach p in the place of n

else

make a cluster for each n_i for $1 \leq i \leq m$, $i \neq q$.

cluster n and n_{q_1} and represent it as a new node r

such that $r = n \cup n_{q_1}$ and

$$n(r) = \max\{n(n_{q_1}) + n(n), n(n_{q_2}) + c(n_{q_2}, n) + n(n)\}$$

delete n and all n_i from the tree

attach r in the place of n

Step 3.5 Go to step 1

We show in the following theorem that Algorithm 3.3.4 produces an optimal allocation in the restricted case.

Theorem 3.3.2 Algorithm 3.3.4 produces an optimal allocation when the following assumptions are satisfied:

- The task precedence graph is a tree.
- There are a greater number of processors

then the number of leaf nodes in the task precedence graph. If no more clustering is required in the sense that the groups of proper nodes have been obtained.

Proof of Theorem 3.2.2 If assumption 2) is satisfied, α_i is always less than or equal to $\alpha_i - \alpha(x)$ in step 3.4 of Algorithm 3.2.4. Thus, the else part of step 3.4 will always be executed. Assumptions 1) and 2) imply that the finishing time of the node x (the time required to complete the node x and its descendants) is only dependent on how the node x and its descendants are allocated. Thus, we only need to show that Algorithm 3.2.4 generates an optimal solution for a subgraph of the given task precedence graph. Suppose that there is a subgraph consisting of a node x and its predecessors α_i , for $1 \leq i \leq m$. Let the allocation produced by Algorithm 3.2.4 be A , and the time required to complete the tasks in this subgraph using A be C_A . Suppose that we have another allocation, called B , for this subgraph, and the time required to complete it in B is C_B , and $C_B < C_A$. In the following, we show that $C_B < C_A$ cannot be true. In A , all α_i except α_p will be allocated to different processors, and in each cycle of Algorithm 3.2.4 only one task is chosen to cluster with its successor, and the addition of any other task to this cluster will not reduce the completion time by the assumption 2). In B a different task, $\alpha_{j'}$, other than α_p , would be selected in step 3.4 to cluster with its successor x . In step 3.4, α_p will be chosen instead of $\alpha_{j'}$, because $\alpha(\alpha_p) + \alpha(\alpha_p, x)$ is the largest.

$$\text{Thus, } C_B = \max\{\alpha(\alpha_{j'}) + \alpha(\alpha_{j'}, x) + \alpha(x), \alpha(\alpha_p) + \alpha(x)\}$$

$$\text{Hence we have that } \alpha(\alpha_p) + \alpha(\alpha_p, x) > \alpha(\alpha_{j'}, x)$$

$$C_B = \alpha(\alpha_{j'}) + \alpha(\alpha_{j'}, x) + \alpha(x) \tag{3}$$

$$\text{In } A, C_A = \max\{\alpha(\alpha_p) + \alpha(\alpha_p, \alpha(\alpha_p)) + \alpha(\alpha_p, x) + \alpha(x)\} \tag{4}$$

Therefore, there are two cases:

case 1) if $\alpha(\alpha_p) + \alpha(x) \geq \alpha(\alpha_{j'}) + \alpha(\alpha_{j'}, x) + \alpha(x)$ then in (3)

$$C_A = c(n_A) + c(x) \quad (3)$$

From (3) and (5), we know that $C_A \geq C_B$ is a contradiction

over (5) if $c(n_A) + c(n_{A,x}) + c(x) > c(n_A) + c(x)$ then in (5)

$$C_A = c(n_A) + c(n_{A,x}) + c(x) \quad (6)$$

From (3) and (6), to satisfy $C_A < C_B$, the following equation should be true:

$$c(n_A) + c(n_{A,x}) + c(x) < c(n_A) + c(n_{A,x}) + c(x)$$

This contradicts the assumption that $c(n_A) + c(n_{A,x})$ is the largest. Thus, no such allocation B exists.

We show the complexity of the Algorithm 3.2.4 in the following theorem.

Theorem 3.2.4 The Algorithm 3.2.4 can run in the time complexity of $O(n)$ where n is the number of nodes in a given DFG.

Proof of Theorem 3.2.4 Step 1 can run $O(1)$ time. Step 2 can also run in $O(1)$ time. Step 3.1 requires $O(1)$ time. Steps 3.3 and 3.2 require visiting ϵ number of nodes, in which ϵ is the number of processors, and thus $O(\epsilon)$ time is required. Step 3.4 also requires visiting ϵ number of nodes, and thus $O(\epsilon)$ time is needed to complete the step. Once they are visited, each node will not be visited again since they are clustered at this step. Thus, the entire algorithm requires visiting at most once each node. Therefore the time complexity of the algorithm is $O(n)$.

3.2.4. Graph Evaluation

Now we extend our graph size determination approach to the general case in which the task precedence relations are represented by a directed graph. We first define depth and height in the directed graph.

Definition 5.3.15 *A depth of a node in a graph is the length of the longest path from the highest ancestor of that node*

A node having no incoming edge is of depth 0

Definition 5.3.16 *A height of a graph is the largest depth of its graph*

A *join graph* G_j and a *task precedence graph* G_p are defined in the same manner as the *join tree* and *task precedence tree* except that they are graphs. The nodes in the join graph are called *join nodes* to distinguish them from the nodes in the task precedence graph. In the graph cases, we also analyse joins to determine the proper join node. We first build a join graph and apply the algorithm 5.3.3 to determine join node. The following is an algorithm to build the join graph.

Algorithm 5.3.3 Build Join Graph

input: A task precedence graph $G_p = (V, E)$, $|V| = n$ and $|E| = m$

output: A join graph in which each node represents a sub-tree in the graph

Step 1 Determine the depth of each node in G_p

Step 2 Set all the nodes "unmarked"

Step 3 For a depth $d = 0$ to a height h do

For each node v of depth d do

Step 3.1 Find predecessors, v_1, v_2, \dots, v_n of v having depth $d - 1$

if they are not all "marked" then

Call Join-Analysis(v_1, v_2, \dots, v_n)

Set v_1, v_2, \dots, v_n "marked"

if the out degree of v is greater than 1 then

set v "marked"

Connect the join node with the existing join nodes

Step 2: Find the successors v_1, v_2, \dots, v_m of v having depth $d + 1$
 Call *Gain-Analyze*(v, v_1, v_2, \dots, v_m)
 Set v_1, v_2, \dots, v_m 'marked'
 Connect the gain node with the existing gain nodes

In Step 1, the depth of each node can be determined using a breadth first search method. Each node may be involved in more than one path, and consequently can have a different length for each path. In such cases, by the definition of depth, the largest value is set to the depth of that node. In Step 2, all the nodes are initially set as 'unmarked', meaning that the node is not involved in any clustering. In Step 3, the gain graph is built by visiting each node. First, a node having smallest depth is chosen, and the predecessors are first checked to determine whether all of them are already included in any cluster. If they were not included in any cluster, then *GAIN* is calculated. The predecessors are all set 'marked', and the node is set 'marked' unless it has only one successor. A gain node representing the nodes under consideration is created. Edges are established from the existing gain nodes to the newly created gain node if there is a node common to both the new gain node and the existing nodes. The same steps are applied to the successor nodes except that all the nodes are 'marked'.

The time complexity of this algorithm is analyzed as follows.

Steps 1 and 2 require visiting each node once, and thus $O(n)$ time is required. The complexity of Step 3 is dominated by the procedure *Gain-Analyze*, and thus each pass of Step 3 requires $O(m + m')$ where m and m' are the number of predecessors and the number of successors, respectively. Step 3 needs to be executed for each node in the graph. Because $m + m'$ is the total number of edges in a node, each edge

needs to be visited at most twice and thus the time complexity of this step is bound to $O(e)$. Therefore, the overall complexity of the algorithm 3.3.3 is bound to $O(e)$.

Once the gain graph is built, we can apply the algorithm 3.3.3 to determine the proper gains. Thus the newly obtained task precedence graph as a result of applying the algorithm 3.3.3 can be used as input for the scheduling stage.

To illustrate our approach we use the example used in [38]. The task precedence graph for the FFT (Fast Fourier Transformation) is shown in Figure 3.9 using the same notation as the former example. Using the algorithm 3.3.3, the gain graph is obtained as shown in Figure 3.10.

We also show a schedule for the FFT problem in the case of four processors in Figure 3.11. McCreary [38] extended her approach by adding analyzing techniques for some special patterns of dependency relations among primitive nodes. Note that such has improved method the former task precedence test example shown in Figure 3.3 results in the same schedule because the former example does not include primitive nodes. Compared to McCreary's approach, our approach has the following advantages: First, our approach is more efficient in terms of time complexity. As shown above, the time complexity of our approach is $O(\max(\alpha \log n, e))$ in which $\alpha \log n$ is for the algorithm 3.3.3 and e is for the algorithm 3.3.4. The time complexity of McCreary's approach is dominated by the parsing algorithm of $O(n^3)$. Second, our approach is more general in the sense that any acyclic task graph can be analyzed. McCreary's method can handle only a few patterns of dependency relations among nodes. Third, our approach considers the execution and communication times as primary factors from the beginning and selects the best candidate from the entire task precedence graph based on the analysis of such times. As discussed before, McCreary's method, however, may lose the clustering opportunity at later stages because the decomposition of the graph is done without using such time information.




Figure 4.8. A task precedence graph for the FFT problem.




Figure 5.15: A game graph for the FTD problem




Figure 5.11 A Schedule for the FFT problem

CHAPTER 4 DISCUSSION

In this dissertation, we have provided (a) the computation model PROOF as a basis for expressing parallelism, (b) the IPR and the transformation rules as a basis for exploiting parallelism and (c) graph size determination techniques for these patterns of parallelism.

We have developed a computation model PROOF in which an object-oriented paradigm is reinterpreted into a functional paradigm without sacrificing the advantages of either. PROOF allows the expression of parallelism on two different levels: object level and method level. The abstracted transparency and history mechanisms are integrated via the pseudo-functions Σ . One of the problems in a parallel object-oriented computation model is to smoothly integrate synchronization constructs and coherence. In PROOF, this integration is achieved by separating the synchronizations that consist of each method from its pure operations.

We have developed an intermediate form IPR to make parallelism explicit in the PROOF/L programs. The IPR is a hybrid graphical representation in which object-level behavior is represented as a Petri net and method-level behavior is represented as a set of linear codes and their data dependency relations. We have also developed transformation rules from the PROOF/L programs to the IPR and proved that the transformation preserves the semantics of the PROOF/L program by showing that the interleaving semantics of the PROOF/L program can be retrieved from the operational semantics of the IPR. The separation of semantics on the two different levels makes verification of the program easy. By introducing the IPR, the semantics

and performance issues of the program can be treated separately. The IFP can serve as a task precedence graph for grain size determination analysis. The programmer's responsibility to ensure correct synchronization and communication is reduced by inserting the proper code during the transformation phase. Considering the fact that coding of synchronization and communication is highly error-prone, this automatic coding reduces errors and thus benefits the software development.

We have presented a two-level allocation approach in which the objects are partitioned into a set of clusters and in each cluster the proper grain size is determined with a each method. We have developed the grain size determination algorithms for different types of partition: pipe-lined, tree and graph. In the case of pipe-lined partition, we show that our algorithm can find optimal solutions. In the case of tree-partition and graph-partition, we compare our approach to the existing methods and show that our approach can perform better in terms of the time complexity.

There are many potential directions for the future work. One immediate direction is to improve the performance of the target code by developing optimization techniques. We have experimented on various PASCAL programs, including factorial problem, bounded buffer problem, dining philosopher problem, warehouse management problem and air base defense simulation problem. We have compared the performance of the target C codes generated by our transformation system with the performance of the target C codes written by hand in terms of the completion time. Because the current transformation system does not include optimization techniques, the performance of the generated code is not as good as the code implemented directly on the target system. For example, we have extensively tested several versions of the factorial program. The speed ups of the two factorial programs, the generated code and the directly written code by hand, are almost the same, but the absolute

completion time of the generated code is at least twice longer than the completion time of the directly written code. When the programs involve extensive manipulation of list data types, the generated code requires much more time than the written code. We believe that it is due to inefficiency of current list handling routines implemented in the front-end transformation. In addition to the code optimization techniques, we need to develop techniques that can reduce unnecessary data movement during object creation. Another interesting aspect is to extend current work to real-time systems, such as process control systems, and to robotics. In particular, its extension to distributed real-time applications can be of immediate interest. To do so, we will need to extend or modify the computation model in order to address real-time specific issues, such as timing constraints and fault-tolerance. In addition, scheduling and allocation strategies need to be changed, since in real-time applications the goal of scheduling and allocation differs so that their goal is not to simply reduce the completion time but to schedule and allocate to meet the deadline of each task. It will also be worthwhile to investigate how we can improve our computation model PROOF to allow dynamic creation of objects. This extension will also require appropriate adjustments in the entire programming system, including adaptation of dynamic scheduling and allocation approaches. We will also incorporate the extension of the computation model with the array constraints so that arrays of objects can be easily created by programmers.

Current transformation rules for the front-end translation are given based on lowest semantics, which is a very weak version of non-strict semantics. In order to make our approach more versatile, we also need to develop different transformation rules for non-strict semantics. One way to detect parallelism in non-strict semantics is to evaluate a function and its arguments at the same time. Then the effect of the non-strict semantics can be determined. In addition, the transformation rules

for exploiting massive parallelism by lazy evaluation need to be developed. The comparison of these three ways of exploiting parallelism will also be of interest.

Current implementation does not fully support the CREW model of execution, although the computation model PRGOL supports such a model. Each object is either considered as a read-only object or a writable object, but objects can be combinations of both. We will investigate a technique to fully support the CREW model, including efficient synchronization mechanisms for supporting parallel execution.

APPENDIX

A1. Syntax for PROOF/L

```

proofl ← 'program' ID " " class_list obj_list body_list "end"
body_list → body_def body_list
      | body_def
body_def → "body" "of" "object" ID " " func
obj_list → obj_list obj_def
      | obj_def
obj_def → "expression" "of" obj_list "is" "expression" "of" ID
      " " obj_list "}"
expression → "actions"
      |
obj_list → obj_list var
      | var
var → LETTERS
      | data_type
class_list → class_list class_def
      | class_def
class_def → "class" ID "{" data_list "}" var_list super_class
      method_def "end" "class"
super_class → "superclass" " " ID "{" data_list "}" "inherit" " "
      inherit_opt "
inherit_opt → "all"
      | id_list
id_list → id_list " " ID
      | ID
var_list → composition var
      |
var → var "B" data
      | data
method_def → method method_def
      | method
method → "method" ID "(" method_in ")" guard_data "expression" func
method_in → input_list "→" output_list
      |
input_list → data_list
data_list → data " " data_list
      | data
data → ID " " data_type
      | ID " " ID
      | ID " " ID " (" data_type ")"
      | ID " " ID " (" ID ")"
      | data_type

```

```

data_type -> 'int'
| 'boolean'
output_list -> output ',' output_list
| output

guard_delta -> 'guard' '[' head_exp ']'
|

func -> 'alpha' ID '[' func_list ']'
| 'beta' '[' func_list '[' ID ']' func_list ']'
| 'gamma' '[' head_exp_list '[' '[' ']' func_list ']'
| 'delta' '[' func_list '[' ID ']'
| 'eta' ID '[' func_list ']'
| func
stat -> 'while' '[' head_exp ',' func ']'
| 'if' '[' head_exp ',' func ',' func ']'
| ID '[' ID '[' average_exp
| exp
average_exp -> func
|
exp -> exp
| 'bexp' exp exp
head_exp -> head_exp exp exp
| 'head' head_exp
| 'True'
| 'False'
head_exp_list -> head_exp ',' head_exp_list
| head_exp
exp -> ID
| 'MINUS'
| 'PLUS'
| '[' func_list '['
| '[' func_list '['
func_list -> func ',' func_list
| func
bexp -> 'lt'
| 'lt='
| 'gt'
| 'gt='
| 'neq'
head_exp -> 'lt'
| 'lt='
| 'gt'
| 'gt='
| 'neq'

```


All PROOF/3 program examples, their IFFs and generated target code²

Bounded Buffer PROOF/3 program³

Program Bounded-Buffer

```
class buffer (datatype, int int)
  comparison
    case int(datatype) & count: int
  method get (buf buffer → buffer, datatype)
    guard (x: buf count < 0)
    expression
      [x]ok, do([buf: store, buf count], handle: handle)
  method put (buf buffer, x: datatype → buffer)
    guard (x: buf count < size)
    expression
      appendLeft(x), set(buf: store, buf count)
  method isEmpty (buf buffer → boolean)
    expression
      = buf count < 0
  method length (buf buffer → int)
    expression
      buf count
end class

class producer (datatype)
  method produce ( → datatype)
    ...
end class

class consumer (datatype)
  method consume (datatype → )
    ...
end class

Object buf: instance of buffer
Active Object producer: instance of producer
Active Object consumer: instance of consumer

Body of object producer:
  while (True, & [ buf ] put(buf, produce){})

Body of object consumer:
  while (True, & [ & [ buf ] consume([get(buf)])]
and
```

²The target code only for the bounded buffer problem is included

³The program is different from the one shown in Chapter 2 since the class *buffer* is embedded into the class definition

An IPN for g_{cat}



An IPN for g_{cat}




```

        Channel(channel, buf where, BUFSIZE * 4,
        buf count = Channel::kbufcnt);
        break;
    }

    case 1: /* File 2 ready */
    {
        int
        Print("Consumer 245",0);
        received[ConsumerID].message_done();
        make_message(message,1,0,0);
        waitToPut[ID].message(buf where);
        make_message(message,1,0,0);
        waitToPut[ID].message(buf count);
        process[ProducerID].message(buf where);
        received[ProducerID].message(buf count);
        break;
    }

    }
    Print("print count = %d",buf count);
    Print("buffer is");
    fflush(stdout);
    for (int i=buf count; i++)
        Print("%d ",buf where[i]);
    Print("\n\n\n\n\n\n\n\n\n\n");
}

}

void _producer(Process *p, Channel *chanP, Channel *reply,Channel
*response[])
{
    /*network buffer 10000 */
    struct buffer {0000};
    0000 message;
    int 10000;
    int 1;
    int val = 1, /* for count */

    p = p;
    message = 0000 + (val % (sizeof(0000)));
    reply = reply;
    do{10000;10000;10000}
    {
        process[Producer, 40000];
        Channel(channelP, val);
        Channel(channel, buf where, BUFSIZE * 4);
        buf count = Channel::kbufcnt;
        put(buf, 10000, 10000); /* changed the 3rd parameter */
    }
}

```

[illegible]

```

{
  if ((Toplevel() = CHANNEL()) == NULL)
    abort();
}
Programc[SOFT_CHANNELS] = NULL;
Toplevel[SOFT_CHANNELS] = NULL;
control = ProcessQueue(S, 2, Programc, Toplevel);
if (control == NULL)
  abort();
is = Primitivef_main(S, 2, Programc, Toplevel);
if (is == NULL)
{
  abort();
}
/*CALL_TRACE(0),*/
FreeProc(control, is, NULL);
}

```

A generated Java C code corresponding to a cluster of neurons

```

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <sys/mman.h>
#include <sys/time.h>
#include <sys/resource.h>

void _connect(Process *p, Channel *input, Channel *output[], Channel *type)
{
    Channel *mych0;
    struct buffer buf0;
    struct buffer buf1;
    struct buffer buf2;
    MSG * message;
    int len0;
    int size; /* the vector of
    int i;
    p = p;

    mych0 = Channel();
    if (mych0 == NULL)
        abort();
    message = (MSG *)malloc(sizeof(MSG));
    size = 1024;
    struct stat = stat();
    for (i=0; i<1024; i++)
    {
        write(message+message, 0, 1, 40);
        read(input[i], message, size);
        read(output[i], message, size);
        getLinkof, 4, 1024, 1024);
        connect(connect, 1024);
        write(message+message, 0, 1, 1024);
        read(output[i], message, 1024-size);
        write(message+message, 0, 1, 40);
        read(output[i], message, 1024-size);
    }

    CheckListInput, 0;
}

void _add(Process *p, Channel *input[], Channel *output[])

```

```

1  PROCESS spt, /% consumer of
   Channel *sptch; /% Synchronisation channel of

   sptch = ChannelGet();
   if (sptch == NULL)
     abort();

   p0 = ForkJoin_consumer(0, 2, sptch, Progress, Toprec);
   if (p0 == NULL)
     abort();

   Finish(p0);

   ChannelPut(sptch);
}

int main()
{
  PROCESS scontrol;
  PROCESS spt;
  Channel *sptch(INPT_CHANNELS+1);
  Channel *sptchrec(INPT_CHANNELS+1);
  int i;
  int, short, action(ACT_WAIT);
  for (i=0; i< INPT_CHANNELS; i++)
  {
    if (CProgress[i] = ChannelGet() == NULL)
      abort();
  }
  for (i=0; i< INPT_CHANNELS; i++)
  {
    if (CToprec[i] = ChannelGet() == NULL)
      abort();
  }
  Progress[INPT_CHANNELS] = NULL;
  Toprec[INPT_CHANNELS] = NULL;
  control = ForkJoin_consumer(0, 2, Progress, Toprec);
  if (control == NULL)
    abort();
  spt = ForkJoin_consumer(0, 2, sptch, sptchrec);
  if (spt == NULL)
    abort();
  Finish(control, spt, NULL);
}

```


A generated Java C code corresponding to the method definitions

```

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/stat.h>

void getStructBuffer(char*, struct buffer *out0, int *out00)
{
    int i=000;
    int i=000;
    int i=000;
    decOffsetForward, &i=000;
    head(buf->data, &i=000);
    i=000 = 0;
    tail(buf->data, i=000);
    out0->store = i=000;
    out0->count = i=000;
    out00 = i=000;
}

void get(struct buffer *inf, int x, struct buffer *out0)
{
    int i=000;
    int i=000;
    int i=000;
    int i=000;
    int i=000;
    int i=000;

    i=000 = buf->data;
    i=000 = buf->count;
    i=000 = i=000;
    incOffset, &i=000;
    i=000 = i=000;
    i=000 = 0;
    append_right0, i=000, i=000;
    out0->data = i=000;
    out0->count = i=000;
}

void produce(struct produce_class *producer, int *out0)
{

```

```

        cout<<endl;
        p1->name = p2->name;
        *p1->name = *p2->name;
    }

void Car::Construct Car(const Car& car, int a)
{
    Car = car;
    a = a;
    *p1->name = *p2->name;
}

```

Encapsulated PROLOG-L program

program fact

class $is(ask)$

method for $x:is(ask) \rightarrow ask$

expression

$is(askable(x) \rightarrow \&ask, is(askable(x) \rightarrow x (has ? \rightarrow a 1)) \&)$

end class

Object Object 01 instance of 01

Body of object 01 : fact 0

end

An IFB for fact



Using Philosophers FORD/L program

```

program dining_philosopher

class chopstick
  composition
    stick : int

  method acquire(a chopstick => chopstick)
    guard { a < stick 0 }
    expression
      inc stick

  method release(a chopstick => chopstick)
    guard { a < stick 1 }
    expression
      dec stick

end class

class philosopher
  composition
    dump : int

  method think(philosopher => int )
    expression
      delay(0)

  method eat( p philosopher => int )
    expression
      delay(0)

end class

Active Object p1 : instance of philosopher
Active Object p2 : instance of philosopher
Active Object p3 : instance of philosopher
Active Object p4 : instance of philosopher
Active Object p5 : instance of philosopher

Object c1 : instance of chopstick
Object c2 : instance of chopstick
Object c3 : instance of chopstick
Object c4 : instance of chopstick
Object c5 : instance of chopstick

```

Body of object p0

```
while( true, (R1(x)) is acquire(x), R2(x)) is acquire(x), set(p0,
    R1(x)) is release(x), R1(x))
is release(x), then(p1)) 0
```

Body of object p0

```
while( true, (R1(x)) is acquire(x), R2(x)) is acquire(x), set(p0,
    R1(x)) is release(x), R1(x))
is release(x), then(p1)) 0
```

Body of object p0

```
while( true, (R1(x)) is acquire(x), R1(x)) is acquire(x), set(p0,
    R1(x)) is release(x), R1(x))
is release(x), then(p1)) 0
```

Body of object p0

```
while( true, (R1(x)) is acquire(x), R1(x)) is acquire(x), set(p0,
    R1(x)) is release(x), R1(x))
is release(x), then(p2)) 0
```

Body of object p0

```
while( true, (R1(x)) is acquire(x), R1(x)) is acquire(x), set(p0,
    R1(x)) is release(x), R1(x))
is release(x), then(p2)) 0
end
```

An IFB for Dining Philosophers
 In the following we present the textual representation of an active object *pl*

Active Object *pl* : *philosopher*

File *l.macro*

Node Number	Node Name	Number of Input	Input Node	Number of Output	Output Node Number
51	DEST	0	29,34	0	31,OUTPUT
52	TRUE	1	34	1	53
54	COPY	1	53	0	55,56
55	MERGE	2	54,56	1	58
57	ϕ	0	INPUT	1	59
60	clock	1	61	1	62
63	LATCH	2	62,64	1	65
67	ASSIGN	1	67	0	61,63
67	release	1	67	1	68
68	LATCH	2	68,67	1	69
68	ASSIGN	1	68	0	69,69
69	release	0	69	1	69
69	LATCH	2	69,69	1	69
69	exit	1	69	1	69
69	LATCH	2	69,69	1	69
69	ASSIGN	1	69	0	61,67
69	acquire	1	69	1	69
69	LATCH	2	69,67	1	69
69	ASSIGN	1	69	0	61,64
69	acquire	1	69	1	69
69	LATCH	2	69,69	1	69

;

REFERENCES

- [1] T. L. Adam, K. M. Chandy, and J. R. Dinkels. A comparison of list scheduling for parallel processing systems. *Communications of the ACM*, 17(12):883-890, 1974.
- [2] G. Agta. *Agata: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, Massachusetts, 1988.
- [3] F. Allen, M. Burke, P. Charles, B. Cyren, and J. Forester. An extension of the PTRAN analysis system for multiprocessing. *Journal of Parallel and Distributed Computing*, 5(4):617-640, October 1987.
- [4] F. Agostini. A parallel object-oriented language. In A. Yonezawa and M. Tokura, editors, *Object-Oriented Concurrent Programming*, pages 189-228. MIT Press, 1987.
- [5] Ardent Computer Co. *Programmer's Guide*, 1989.
- [6] Z. Aruda and Arvind. P-ENC: A parallel-intermediate language. In *Proc. of the Functional Programming Languages and Computer Architecture*, pages 359-365, 1989.
- [7] Arvind and K. Elmasri. Future-oriented programming on parallel machines. *Journal of Parallel and Distributed Computing*, 5(3):466-493, Oct. 1988.
- [8] J. Backus. Can programming be liberated from the von Neumann style? A functional style and its algebra of programs. *Communications of the ACM*, 21(6):632-641, Aug. 1978.
- [9] H. E. Bal, Steven J. G., and A. S. Tanenbaum. Programming languages for distributed computing systems. *ACM Computing Surveys*, 23(2):353-399, 1990.
- [10] H. P. Barendregt, M. C. J. D. van Eekelen, J. R. W. Glauert, M. J. Flanagan, and M. B. Sleep. Towards an intermediate language based on graph rewriting. In *Proc. of the FASES Conference*, LNCS 648, pages 156-176, 1991.
- [11] J. Boyan, M. Deneys, and D. Yungster. The GFL2 experiments. In *Proc. of the 1988 Annual International Symposium on Computer Architecture*, pages 106-116, 1988.
- [12] A. Black, N. Elmqvist, E. Jø, and E. Levy. Object structure in the Euler-oid system. In *Proc. of Object-Oriented Programming Systems, Languages and Applications*. ACM SIGPLAN Notices 23(11), pages 76-86, 1988.

- [12] P. Denck-Hansen. The programming language concurrent Pascal. *IEEE Transactions on Software Engineering*, 1(2):189-202, June 1975.
- [13] L. DeRemeter. *Programming Expert Systems in OPS3*. Addison-Wesley, Reading, Massachusetts, 1982.
- [14] R. Curran and B. Colander. Loads in context. *Communications of the ACM*, 25(4):444-458, 1982.
- [15] C. Y. Chin and E. Eising. Task scheduling algorithms for multiprocessors and database computers. *IEEE Trans. on Computers*, C-33(1):99-108, 1984.
- [16] W. W. Chu, L. J. Hallaway, M. T. Lee, and E. Eln. Task allocation in distributed data processing. *IEEE Computer*, 13(11):37-48, 1980.
- [17] A. Church and J. B. Boyer. Some properties of combinators. *Transactions of American Mathematical Society*, 28:419-452, 1958.
- [18] E. L. Clark and S. Gregory. PARLOG: Parallel programming in logic. *ACM Transactions on Programming Languages and Systems*, 6(1):1-49, 1984.
- [19] Gnu Research Inc. *System (GFT) Software Manual*, 1984.
- [20] A. L. Davis and R. M. Keller. Data flow program graphs. *Computer*, 15(2):56-61, 1982.
- [21] Department of Defense. *Architecture Manual for the Ada Programming Language*, 1983. ARJ/MBL-STD-1815A-1003, 1983.
- [22] E. W. Dijkstra. Guarded Commands, nondeterminacy and formal derivation of programs. *Communications of the ACM*, 18:978-991, 1975.
- [23] E. Eln. Recursive models of task assignment scheduling in distributed systems. *IEEE Computer*, 15(1):50-58, 1982.
- [24] G. J. El Jaiouli and W. E. Ryan. Distributed computation on network computers. *IEEE Transactions on Computers*, C-39(2):418-425, 1990.
- [25] E. P. Evanson, J. N. Gray, R. A. Loebe, and L. L. Traiger. The solution of consistency and predicate locks in a database system. *Communications of the ACM*, 19(11):626-633, Nov 1976.
- [26] S. Edelkamp. *NETS: A System for Expressing and Using Real World Knowledge*. MIT Press, Cambridge, Massachusetts, 1978.
- [27] J. R. Healey. The Flex/32 for real time multiprocessor simulation. In W. J. Karolus, editor, *Multiprocessors and Array Processors*, pages 121-126. Santa Lisa Concourse, Inc., 1987.
- [28] M. J. Heule. Some computer organizations and their effectiveness. *IEEE Transactions on Computers*, C-33(2):548-565, 1973.
- [29] M. J. Heule. Parallelism in random access machines. In *Proc. 1986 ACM STOC*, pages 118-119, 1986.

- [21] J. Foster. *Shared Flow Concepts in Parallel Programming*. Prentice Hall, Englewood Cliffs, NJ, 1980.
- [22] M. H. Gandy and D. S. Johnson. *Computers and Integrability: A Guide to the Theory of NP-completeness*. Prentice, San Francisco, CA, 1979.
- [23] J.-L. Gaudel and L.-T. Lee. Decompose: A methodology for programming multiprocessor systems. *Journal of Parallel and Distributed Computing*, 196: 119, 1993.
- [24] B. Gobleberg. Multiprocessor execution of functional programs. *Journal of Parallel Programming*, 17(2): 425-433, 1989.
- [25] J. Graham and J. Reagin. Expert compilation on the CDC microvax computer. In W. J. Karplus, editor, *Multiprocessors and Array processors*, pages 147-158. Simulation Councils, Inc, 1983.
- [26] A. S. Greenberg and J. W. S. Liu. Mosaic: An object oriented macro data flow system. In *Proc. of Object Oriented Programming Systems, Languages and Applications 1987*. ACM SIGPLAN Notices 22(17), pages 31-42, 1987.
- [27] V. B. Gryn and J. A. Edwards. Optimal partitioning of workload for distributed systems. In *Object of Papers, COMPCON Fall 86*, pages 352-357, 1986.
- [28] S. H. J. Halstead. Multilisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Applications*, 7(4): 504-538, 1985.
- [29] W. D. Hillis. *The Connection Machine*. The MIT Press, Cambridge, Massachusetts, 1985.
- [30] C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8): 666-677, Aug 1978.
- [31] B. Hockney and C. Jesshope. *Parallel Computers*. Adam Hilger, 1980.
- [32] P. Hudak. Pure functional programming. *IEEE Computer*, 19(3): 60-76, 1986.
- [33] C. Huxson, T. Macho, J. Davies, M. White, and B. Leasure. The KAP-205: An advanced source-to-source translator for the Cyber 205 supercomputer. In *Proc. of International Conference on Parallel Processing*, pages 517-523, 1988.
- [34] J.-J. Huang, F. D. Anger, Y. C. Chow, and C. Y. Lee. Scheduling precedence graphs in systems with interprocessor communication buses. *IEEE Journal on Computing*, 15(2): 244-253, 1989.
- [35] R. Huang. Advanced parallel processing with supercomputer architecture. *Proceedings of the IEEE*, 76(10): 1245-1279, 1987.
- [36] Isaac Idd. *Green Programming Manual*, 1986.
- [37] R. Jagannathan, A. N. Grewing, W. T. Ziegler, and R. E. S. Lee. Dataflow-based methodology for source-to-parallel compilation on a network of workstations. In *International Conference on Parallel Processing*, volume 3, pages 269-285, 1983.

- [15] R. Jagnathan and A. A. Pessen: The GLT programming language. Technical Report 583 CSL-80-11, 350 International Computer Science Laboratory, 1980.
- [16] D. C. Kahan and K. H. Lee: Interference in actor based concurrent object-oriented languages. In *Proc. of the 2nd European Conference on Object-Oriented Programming*, pages 133-145, 1989.
- [17] E. Kumbhar and M. Senechal: Practical multiprocessor scheduling algorithms for efficient parallel processing. *ACM Transactions on Computers*, C-32(11):1003-1009, 1984.
- [18] R. M. Keller: Formal verification of parallel programs. *Communications of the ACM*, 15:471-484, 1972.
- [19] S. J. Kim: A general approach to multiprocessor scheduling. Technical report, University of Texas at Austin, 1988.
- [20] H. T. King: Notes on VLSI compilation. In D. J. Evans, editor, *Parallel Processing Systems*. Cambridge University Press, 1977.
- [21] R. Liberman: Concurrent object-oriented programming in Act 1. In A. Norman and M. Tokoro, editors, *Object-Oriented Concurrent Programming*, pages 3-36. MIT Press, 1992.
- [22] Y.-M. Lu: Heuristic algorithms for task assignment in distributed systems. In *Proc. 4th Int'l Conf. on Distributed Computing Systems*, pages 59-65, 1984.
- [23] Ibmco Ltd: *The Transputer Databook*. Ibmco Ltd., Bristol, UK, 1989.
- [24] F. R. Ma and E. Y. S. Lee: A task allocation model for distributed computing systems. *ACM Transactions on Computers*, C-31(1):41-47, 1982.
- [25] C. McGarry and B. Gall: Automatic determination of grain size for efficient parallel processing. *Communications of the ACM*, 35(9):1053-1074, 1992.
- [26] C. McGarry and B. Gall: Efficient exploitation of concurrency using graph decomposition. In *International Conference on Parallel Processing*, pages 166-183, 1990.
- [27] J. McLaw: SP5AL: macros and libraries in a single environment language. *Language reference manual, version 1.2*. Technical Report Technical Report 85-140, LLNL, 1985.
- [28] G. M. Moryhead: Active objects in hybrid. In *Proc. of Object-Oriented Programming Systems, Languages and Applications*, pages 242-253, 1991.
- [29] R. S. Nikhil, K. Pingali, and Arvind: M. Newman. Technical Report Computer Science Group Memo 85, Laboratory for Computer Science, MIT, 1986.
- [30] C. E. Patschke: The serializability of concurrent database updates. *Journal of ACM*, 20(4):621-653, Oct 1979.

- [16] J. F. Palmer. The MCMC family of parallel supercomputers. In W. J. Karplus, editor, *Multiprocessors and Array processors*. Simulation Councils, Inc, 1981.
- [17] J. L. Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice-Hall, Englewood Cliffs, New Jersey, 1981.
- [18] G. F. Plett, W. C. Bransley, D. A. George, S. L. Harvey, E. P. McCallife, and E. A. Miller. The data oriented parallel processor prototype (DOP). In *Proc. of International Conference on Parallel Processing*, pages 261-711, 1980.
- [19] C. D. Polychronopoulos and U. Banerjee. Processor allocation for hierarchical and vertical parallelism and related scheduling bounds. *IEEE Transactions on Computers*, C-30(4):410-426, 1981.
- [20] C. C. Price and S. Kratochvil. Software allocation models for distributed computing systems. In *International Conference on Distributed Computing Systems*, pages 43-48, 1984.
- [21] C. Y. Ramamurthy, J. Srinivasan, and W. T. Tsai. Clustering techniques for large distributed systems. In *IEEE INFOCOM 86*, pages 85-99, 1986.
- [22] Y. Sakai and Y. Homma. Partitioning parallel programs for mainframe. In *ACM Conference on Log and Functional Programming*, pages 200-211, August 1988.
- [23] R. Shapiro. *Concurrent Programming: A progress report*. *IEEE Computer*, 10(3):44-54, 1977.
- [24] C. C. Shen and W. T. Tsai. A graph matching approach to optimal task assignment in distributed computing systems using a heuristic algorithm. *IEEE Transactions on Computers*, 30(2):157-165, 1981.
- [25] S. K. Shukla and M. L. Wolfson. Data flow graph optimization - IF1. In *Proc. of Functional Programming Languages and Computer Architecture Conference*, pages 12-34, 1981.
- [26] R. S. Stone. Multiprocessor scheduling with the aid of network flow algorithms. *IEEE Transactions on Software Engineering*, SE-8(1):55-65, 1981.
- [27] R. S. Stone and R. H. Balkman. Control of distributed processes. *IEEE Computers*, 11(7):57-66, 1978.
- [28] J. Tasi, M. Myerowitz, and R. C. Smith. The Alliant FX/Sequoia: Automatic partitioning on a multiprocessor real supercomputer. In W. J. Karplus, editor, *Multiprocessors and Array processors*, pages 25-64. Simulation Councils, Inc, 1981.
- [29] S. Thomas, P. Gifford, and G. Finkel. Balcon: A shared memory multiprocessor. In *Proc. of the Second International Conference on Supercomputing*, pages 25-36, 1987.
- [30] C. Timmerman and V. Singh. Inheritance and specialization with restricted sets. In *Proc. of Object Oriented Programming Systems, Languages and Applications*, pages 300-312, 1988.

- [79] K. E. Dush. Compilation as partitioning: A new approach to compiling non-strict functional languages. In *Proc. of the Functional Programming Languages and Computer Architecture*, pages 75-88, September 1989.
- [80] J. L. Ullrey. Mathematical programming approaches to system partitioning. *IEEE Transactions on Systems, Man, Cybernetics*, SMC-9:640-648, 1979.
- [81] J. D. Wilson. NP complete scheduling problems. *Journal of Comput. Syst. Sci.*, 14:556-593, 1975.
- [82] S. S. Yen, D. H. Ben, and H. Choudhary. A framework for software development for distributed parallel computing systems. In *Proc. of the Third Workshop on the Future Trends of Distributed Computing Systems*, April 1990.
- [83] S. S. Yen, K. Jin, and D. H. Ben. Trends in software design for distributed computing systems. In *Proc. of the Second Workshop on the Future Trends of Distributed Computing Systems*, pages 144-149, 1989.
- [84] S. S. Yen, K. Jin, and D. H. Ben. FRODO: A parallel object oriented functional computation model. *Journal of Parallel and Distributed Computing*, 13(3):203-213, July 1994.
- [85] S. S. Yen, K. Jin, and D. H. Ben. Software design methods for distributed computing systems. *Journal of Computer Communications*, 1993.
- [86] S. S. Yen, K. Jin, D. H. Ben, M. Choudhary, and G. Qi. An object-oriented approach to software development for parallel processing systems. In *Proc. International Computer Software and Applications Conference (COMPSAC'93)*, pages 403-408, 1993.
- [87] S. S. Yen and C. S. Lu. A structured inheritance network representation for object oriented software design. In *Proc. of 12th Annual Int'l Computer Software & Applications Conference (COMPSAC'93)*, pages 324-330, 1993.
- [88] Y. Yokota and H. Tokoro. Concurrent programming in concurrent Haskell. In A. Yonezawa and M. Tokoro, editors, *Object-Oriented Concurrent Programming*, pages 128-154. MIT Press, 1991.
- [89] A. Yonezawa, J.-P. Boud, and E. Schreyer. Object-oriented concurrent programming in ABCAL. In *Proc. of Object-Oriented Programming Systems, Concepts and Applications: ACM SIGPLAN Notices 1993*, pages 224-235, 1993.

BIOGRAPHICAL SKETCH

Doa-Heun Do was born on December 6, 1957, in Seoul, Korea. He received the B.S. and M.S. degrees in mechanical engineering from the Seoul National University, Seoul, Korea, in 1980 and 1982, respectively. He completed his military duty in May 1983, and worked as a Researcher at Korea Advanced Institute of Science and Technology until May 1984. He received the M.S. degree in computer science from the University of Wisconsin-Milwaukee in 1984. He studied in the Computer Science Program at the Northwestern University from September 1987 to August 1988. He is currently a Ph.D. candidate in the Computer and Information Sciences Department, University of Florida, Gainesville.

I certify that I have read this study and that, in my opinion, it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.



Stephen S. Yau, Yau
Professor of Computer and
Information Sciences

I certify that I have read this study and that, in my opinion, it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.



Yuen Heng Lee
Associate Professor of Computer
and Information Sciences

I certify that I have read this study and that, in my opinion, it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.



Paul W. Chan
Professor of Biochemistry
and Molecular Microbiology

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.



Paul S. Falsworth
Associate Professor of Computer
and Information Sciences

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.



Tim Davis
Assistant Professor of Computer
and Information Sciences

This dissertation was submitted to the Graduate Faculty of the College of Engineering and to the Graduate School and was accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

May 1992



Michael M. Phillips
Dean, College of Engineering

Michelle M. Lockhart
Dean, Graduate School